Proceedings of the Second International Workshop on Nonmonotonic Reasoning, M. Reinfrank et al. (eds.), Springer-Verlag (1989), pp. 220-237

Logics for Inheritance Theory

Richmond H. Thomason
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
U.S.A.

John F. Horty
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
U.S.A.

1. Introduction

In contemporary A.I., there is much room for variation in the distance between theory and implementation. Some implementations (*e.g.,* some parsing algorithms), can actually be shown to compute the relations that are specified by a theory. But sometimes—even though there are theories, and the theories clearly help in a way to inspire and understand implementations—the relationship between the two is looser and, on the whole, more problematic. And sometimes there are no theories at all.

Knowledge representation seems to be a case in which the gap between theory and implementations is yet to be closed. There are now a number of well developed logic-based theories of nonmonotonic reasoning,[1] and numerous frame-based implementations that employ nonmonotonic reasoning in A.I. applications. But it remains to be shown that one of these theories of nonmonotonic reasoning provides a faithful account of a reasonable frame-based system. And even if we consider mathematical models of inheritance, such as those of [24] and [10], which dispense with many of the complexities of actual implementations, there is usually no straightforward way of applying the logical theories to these models. Thus, it is difficult at present to put these theories to work, either in criticizing existing technologies or in designing improvements.

We have no simple cure for the gap, because we believe that it is due in part to the intrinsic difficulty of the research problems. In this paper we propose a diagnosis of the problem; describe the methodology for attacking it that has emerged in a project at Carnegie Mellon University and The University of Pittsburgh; we give a generalization of the difficulties involved in relating logic to nonmonotonic inheritance in semantic nets; and, to illustrate the sort of result that we believe is required, we present a completeness theorem for a very simple case.

2. The theory-implementation gap in knowledge representation

We suspect that existing logical theories of nonmonotonic reasoning have been too conservative, too ready to respect details of logical theories that were developed for entirely different purposes. Familiar logical theories, including intuitionistic, modal, and higher-order logics, were inspired by

[1] We count a logic as well-developed for this purpose only if it has beeg given a model theoretic interpretation.

the need to constr
tions must be mac
runs deep in thes
believe, to provid
nonmonotonic "IS

Our own resea
relatively abstrac
seek to develop m
[24] serves as a r
deals with object
theoretical pursui
it is not the same
important step in

Logical theory
require logics tha
plicable logics do
are generated by
may be unfamilia
techniques in the
objects of logical

This approac
the relation betw
efficiently. (2) T
relation between
of statements fro
than that of the
characterizes the
and complete wit
for the logical lar

As an exercise
with respect to a
noteworthy conse
classical, two-valu
believe, that a n
lattice of truth v

In this paper,
logic, and that t
one. This logic i
inheritance in th

the need to construct theories of mathematical reasoning. And in mathematical reasoning, assumptions must be made explicit, and theorems must be reusable in proving later results. Monotonicity runs deep in these theories, and incremental modifications to standard logics are not likely, we believe, to provide tractable characterizations even of apparently simple applied concepts such as nonmonotonic "IS-A."

Our own research strategy can be summarized as a "bottom-up" approach that seeks to link relatively abstract logical theories to implementations by developing multiple levels of theory. We seek to develop mathematical accounts of inheritance that are relatively close to implementations; [24] serves as a model for this work. It is important to realize that inheritance theory, which deals with objects like graphs and paths rather than with formulas and proofs, is an independent theoretical pursuit with its own methods, intuitions, and results. In some ways it is like logic. But it is not the same as logic. The emergence of inheritance theory as a separate area of inquiry is an important step in relating theories to actual knowledge representation technology.

Logical theory is another necessary part of the picture; but we feel that nonmonotonicity may require logics that are relatively unfamiliar, and so will not be found in the logical literature. If applicable logics do not exist, they will have to be invented; thus, we hope to appeal to intuitions that are generated by our work in inheritance theory in developing suitable logics. Though these logics may be unfamiliar, of course they will have to be rigorously constructed, using metamathematical techniques in the logical repertoire; and (hopefully) they should in themselves be interesting as objects of logical study.

This approach breaks down into the following detailed tasks. (1) To develop a theory of the relation between a net Γ and a consequence A, showing that this relation can be computed efficiently. (2) To develop a model theory, which yields a definition of logical consequence as a relation between a set of formulas Γ and a formula A. We can regard a semantic net as a set of statements from this logic, though the expressive power of the logical language may be richer than that of the nets, allowing other statements to be formed. (3) Thus, to show that the logic characterizes the account of inheritance, we wish to prove that the inheritance relation is sound and complete with respect to logical consequence. (4) We may also wish to devise a proof theory for the logical language that is sound and complete with respect to logical consequence.

As an exercise in applying this technique, we have carried out all of the above theoretical steps with respect to a system of *monotonic inheritance;* the results are described in [20] and [21]. A noteworthy consequence of this project is that the logic even of monotonic inheritance is not the classical, two-valued one. Instead, it corresponds to a well-known four-valued logic. This shows, we believe, that a *nonmonotonic* logic that meets our goals must be based either on this four-valued lattice of truth values or some extension of it.

In this paper, we will show how a version of autoepistemic logic can be based on this four-valued logic, and that this yields a nonmonotonic logic that in many ways is simpler than the modal one. This logic in itself, however, does not provide a theory that is applicable to nonmonotonic inheritance in the direct way we would like; an appropriate theory of *nonmonotonic quantification* needs to be added to it. We will also discuss this problem, and present a result which relates the logical theory to a simplified version of the "credulous" direct semantics for inheritance that is discussed in [25].

3. A four-valued nonmonotonic propositional logic

The language of the propositional logic has disjunction, conjunction, negation, an operator $\triangle$ which in some ways is like the necessity operator of a modal logic, and an operator $\bigtriangledown$ which in in some ways is like the possibility operator of a modal logic. Complex formulas are constructed according to the following rules. (A *literal* is either an atomic formula or $\neg A$, where A is an atomic formula.)

1. If A, B are formulas, so are $\neg A, A \vee B$, and $A \wedge B$.
2. If A is a formula, so are $\triangle A$ and $\bigtriangledown A$.

Models are characterized as follows.

Truth values are subsets of $\{\mathrm{T}, \mathrm{F}\}$. A model $\mathcal{M}$ assigns truth values to atomic formulas anyhow, and truth values to complex formulas in a way that depends on the values of simpler formulas. $\mathcal{M} \models_{\mathrm{T}} A$ means that T belongs to the value that $\mathcal{M}$ assigns to A; $\mathcal{M} \models_{\mathrm{F}} A$ means that F belongs to this value.

The rules of satisfaction are as follows.

$\mathcal{M} \models_{\mathrm{T}} \neg A$ if $\mathcal{M} \models_{\mathrm{F}} A$	$\mathcal{M} \models_{\mathrm{T}} \triangle A$ if $\mathcal{M} \models_{\mathrm{T}} A$
$\mathcal{M} \models_{\mathrm{F}} \neg A$ if $\mathcal{M} \models_{\mathrm{T}} A$	$\mathcal{M} \models_{\mathrm{F}} \triangle A$ if $\mathcal{M} \not\models_{\mathrm{T}} A$
$\mathcal{M} \models_{\mathrm{T}} A \wedge B$ if $\mathcal{M} \models_{\mathrm{T}} A$ and $\mathcal{M} \models_{\mathrm{T}} B$	$\mathcal{M} \models_{\mathrm{T}} \bigtriangledown A$ if $\mathcal{M} \not\models_{\mathrm{F}} A$
$\mathcal{M} \models_{\mathrm{F}} A \wedge B$ if $\mathcal{M} \models_{\mathrm{F}} A$ or $\mathcal{M} \models_{\mathrm{F}} B$	$\mathcal{M} \models_{\mathrm{F}} \bigtriangledown A$ if $\mathcal{M} \models_{\mathrm{F}} A$
$\mathcal{M} \models_{\mathrm{T}} A \vee B$ if $\mathcal{M} \models_{\mathrm{T}} A$ or $\mathcal{M} \models_{\mathrm{T}} B$	
$\mathcal{M} \models_{\mathrm{F}} A \vee B$ if $\mathcal{M} \models_{\mathrm{F}} A$ and $\mathcal{M} \models_{\mathrm{F}} B$	

So far, there is nothing nonmonotonic about this logic. Nonmonotonicity enters with the definition of the implication relation $\mid\!\!\mid\!\!\sim$ between a set Γ and a formula A. The trick will be this: to characterize a *normal* model of a set Γ, and to say that $\Gamma \mid\!\!\mid\!\!\sim A$ iff $\mathcal{M} \models_{\mathrm{T}} A$ for every normal model $\mathcal{M}$ of Γ.

Our first characterization of a normal model of Γ is that it must be informationally minimal. $\mathcal{M}_1 \leq \mathcal{M}_2$ iff for all atomic formulas A, if $\mathcal{M}_1 \models_{\mathrm{T}} A$ then $\mathcal{M}_2 \models_{\mathrm{T}} A$, and if $\mathcal{M}_1 \models_{\mathrm{F}} A$ then $\mathcal{M}_2 \models_{\mathrm{F}} A$. $\mathcal{M}$ is a normal model of Γ iff $\mathcal{M}$ is a model of Γ, and if $\mathcal{M}'$ is a model of Γ and $\mathcal{M}' \leq \mathcal{M}$ then $\mathcal{M}' = \mathcal{M}$.

The resulting logic, N4,[2] is nonmonotonic. For instance, we have

[2] The '4' in this name of the logic is a reminder that it is four-valued. The 'N' is for 'nonmonotonic'. This name is provisional; if the space of similar logics proliferates, it may be necessary to choose a more complicated name. Any resemblance to 'S4' is purely accidental.

$$\Vdash\!\!\sim \nabla p \qquad \text{and} \qquad p \vee [p \wedge q] \Vdash\!\!\sim \nabla\neg q,$$

but

$$\neg p \not\Vdash\!\!\sim \nabla p \qquad \text{and} \qquad q,\ p \vee [p \wedge q] \not\Vdash\!\!\sim \nabla\neg q.$$

It would be misleading to claim that the logic of $\triangle$ could serve as a satisfactory general-purpose epistemic necessity. The most striking difference between $\triangle$ and any modal $\Box$ is the logical equivalence of $\triangle[A \vee B]$ and $\triangle A \vee \triangle B$; such an equivalence certainly doesn't hold for knowledge in general. We can know *that* A or B without knowing *whether* A or B.

However, it has been argued[3] that the four-valued logic on which **N4** is based characterizes the logic of what is known by simple information retrieval systems, and in certain restricted applications (*e.g.,* closed world systems, and many types of semantic nets as well), the equivalence of $\triangle[A \vee B]$ and $\triangle A \vee \triangle B$ may not be so unreasonable. At any rate, as we will show, **N4**—which in many ways is simpler than autoepistemic logic—does seem to provide a promising basis for interpreting monmonotonicity in semantic nets.

We will not develop the logical metatheory of **N4** here; but will leave that task for another paper.[4]

4. Quantification and generic statements

The propositional logic **N4** provides a foundation for interpreting generic statements (statements corresponding to generic IS-A links in nets, of the form $p \rightarrow q$). But it is by no means trivial to extend the propositional logic in the required way. Although generic IS-A links resemble true universal quantifiers in that they are used to make general statements, the analogy does not, as far as we can see, support a definition of $p \rightarrow q$ in terms of a universal quantifier. In fact, the view that nonmonotonic IS-A links have a logical form that somehow involves universal quantification is (like the view that monotonic inheritance corresponds to two-valued logic) an idea that is attractive because of its appeal to the familiar, but that is unsupported by theoretical results.

The following example illustrates the difficulty. Suppose that we extend the logic **N4** by adding a two-place quantifier $\forall$, allowing formulas such as $\forall x(px, qx)$. This formula is true in a model in case for every individual d, if p is true of d then so is q. (The falsity conditions of the formula are unimportant for the purposes at hand.) We then define $p \rightarrow q$ and $p \not\rightarrow q$ as follows.

$$\begin{aligned} p \rightarrow q \quad &=_{df} \quad \forall x([px \wedge \nabla qx], qx) \\ p \not\rightarrow q \quad &=_{df} \quad \forall x([px \wedge \nabla\neg qx], \neg qx) \end{aligned}$$

The theory we have stated yields a relation $\Vdash\!\!\sim$ of logical consequence; but this relation does not correspond to any reasonable inheritance consequence relation, even in the most simple cases.

---

[3]See [1].

[4]**N4** is decidable, and can be axiomatized using a Gentzen *Sequenzenkalkül* without the rule of weakening. This technique for axiomatizing nonmonotonic logics is discussed in [22].

For instance, let Γ_1 be the set $\{pa, p \rightarrow q\}$. This corresponds to the following simple net.

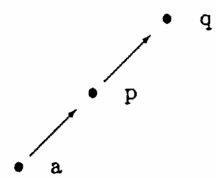

Γ_1

Clearly, qa holds in this net. But—equally clearly—there is a minimal model of Γ_1 in which $\neg qa$ holds. This model satisfies $p \rightarrow q$ because it fails to satisfy ∇qa.

Another kind of difficulty has to do with *preemption*. To see the problem, let Γ_2 be $\{pa, p \rightarrow q, q \rightarrow r, p \not\rightarrow r\}$.

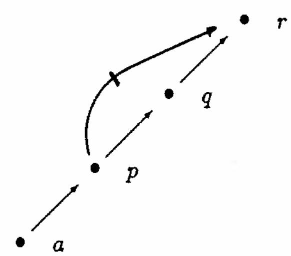

Γ_2

Nothing rules out the model of Γ_2 in which pa, qa, and ra are all true. But this model does not comply with intuitions about inheritance.[5]

The solution to these problems involves a revision of the fundamental notion of a normal model. Simply minimizing information does not suffice to eliminate the abnormality we have identified above, because the inheritance definition of a semantic net is a kind of compromise between concluding too much on the one hand (reaching conclusions that are unwarranted), and, on the other hand, concluding too little (failing to make the best use of general information that is contained in generic links).

Our revised account of normality is as follows: relative to a set of formulas Γ we associate with each model $\mathcal{M}$ three sets, which together serve as measures of normality:

[5]See [24], [11], and [10].

1. A set $inf(\mathcal{M})$ of literals, serving as a measure of the information content of $\mathcal{M}$;
2. A set $inc(\mathcal{M})$ of atomic formulas, serving as a measure of the inconsistencies introduced by $\mathcal{M}$;
3. A set $ano_\Gamma(\mathcal{M})$ of pairs of generic formulas and individuals, serving as a measure of the anomalies introduced by $\mathcal{M}$, relative to the generalizations asserted in Γ;

Given an assignment of normality measures $inf(\mathcal{M})$, $inc(\mathcal{M})$, and $ano_\Gamma(\mathcal{M})$ to $\mathcal{M}$, define a relation $\leq$ of relative normality for models, as follows.

$\mathcal{M}_1 \leq \mathcal{M}_2$ iff either:

1. $inc(\mathcal{M}_1) \subset inc(\mathcal{M})_2$; or
2. $inc(\mathcal{M})_1 = inc(\mathcal{M}_2)$ and $ano_\Gamma(\mathcal{M}_1) \subset ano_\Gamma(\mathcal{M}_2)$; or
3. $inc(\mathcal{M})_1 = inc(\mathcal{M}_2)$ and $ano_\Gamma(\mathcal{M}_1) = ano_\Gamma(\mathcal{M}_2)$ and $inf(\mathcal{M})_1 \subseteq inf(\mathcal{M}_2)$.

The prioritization implicit in this ordering reflects, we believe, the policies that are incorporated in inheritance definitions. First, do not allow any inconsistencies that are not absolutely forced by the hypotheses. Second, maximize use of generalizations in the hypotheses, insofar as this does not conflict with the first priority. Third, do not reach conclusions except insofar as are forced by the hypotheses and the first two priorities.

There are simple intrinsic definitions of $inf(\mathcal{M})$ and of $inc(\mathcal{M})$. The former is the set of literals made true by $\mathcal{M}$, the latter is the set of atomic formulas A such that both A and $\neg A$ are made true by $\mathcal{M}$. The definition of anomaly is not so simple; there seem to be many reasonable ways of characterizing the anomalies of a model with respect to a set of hypotheses. The simplest of these says merely that $<p \rightarrow q, \mathrm{d}>$ is an anomaly if either (1) $p \rightarrow q \in \Gamma$, $\mathrm{d} \in \mathrm{D}_\mathcal{M}$, (where $\mathrm{D}_\mathcal{M}$ is the domain of the model $\mathcal{M}$) and $\mathcal{M} \models_{\mathrm{T,d}} px \wedge \nabla\neg qx$, or (2) $p \not\rightarrow q \in \Gamma$, $\mathrm{d} \in \mathrm{D}_\mathcal{M}$, and $\mathcal{M} \models_{\mathrm{T,d}} px \wedge \nabla qx$.[6]

5. An exercise in modeling inheritance

5.1. Conflict resolution and networks

Networks that can permit positive and negative reasons for the same conclusion, and that are capable of *credulous* reasoning (see [10]), must appeal to some form of conflict resolution—some way of deciding which reasons to accept in case of conflict. In general, this resolution method may involve decisions that are arbitrary; this is the source of multiple extensions that are equally

[6] The subscript 'd' on '$\models$', of course, is keyed to the free variable x.

reasonable, in inheritance systems such as that of [24]. We begin by introducing a framework for discussing conflict resolution.

A network over a set I of individuals and set K of kinds is a pair of subsets of $(\mathrm{I}\times\mathrm{K})\cup(\mathrm{K}\times\mathrm{K})$, the set of *positive links* and the set of *negative links*. We use upper-case Greek letters for networks, 'a', 'b', and 'c' for individuals, 'p', 'q', 'r' for kinds, and arrows '$\rightarrow$' and '$\not\rightarrow$' for links; thus, we may say that $a\rightarrow p\in\Gamma$ and $a\not\rightarrow p\in\Gamma$. In some contexts we use propositional notation, saying that $pa\in\Gamma$ or that $\neg pa\in\Gamma$. But pa and $\neg pa$ are the same things as $a\rightarrow p$ and $a\not\rightarrow p$.

To keep things simple, we'll confine conflict resolution to specific propositions of the form $a\rightarrow p$ and $a\not\rightarrow p$, ignoring generic propositions of the form $p\rightarrow q$ and $p\not\rightarrow q$. So we can think of a resolution function as a partial function f from $I\times K$ to $\{\mathrm{T},\mathrm{F}\}$. Such functions are partial because they need only resolve questions that arise; a question will not arise unless there are active, defeasible pros and cons associated with it. Though the parentheses in the notation '$f(a,p)$' are just delimiters for functional application, it's useful to think of (a,p) as an object—the *question* (a,p) represents an issue as to whether $a\rightarrow p$ or $a\not\rightarrow p$ holds.

As in [10], a *generalized path* in Γ is a sequence of links, positive and negative, belonging to Γ. Γ is *acyclic* if it contains no generalized paths that are acylic. The degree $deg_\Gamma(\sigma)$ of a path σ in Γ (where a path in Γ is a generalized path which, if it has any negative links, has just one such link, in final position) is the length of the longest generalized path in Γ from the head to the tail of σ. The degree of a question (a,p) is the degree $deg_\Gamma(\sigma)$ of an arbitrary generalized path σ in Γ from a to p, if such a path exists; otherwise, $deg_\Gamma(a,p)=0$.

5.2. A crude sort of inheritance

By a simultaneous induction on $deg_\Gamma(a,p)$, we define the sets $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p)$ and $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p)$ of *pros* and *cons* for the question (a,p) (relative to Γ and the resolution function f), and the permission relation $\vdash$ between $<\Gamma,f>$ and paths in Γ. The definition assumes that f is *adequate* for Γ, in a sense to be explained later.

Definition 1

1. *If* $pa\in\Gamma$ *then* $a\rightarrow p\in\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p)$ *and* $<\Gamma,f>\vdash a\rightarrow p$*;*
2. *If* $\neg pa\in\Gamma$ *then* $a\not\rightarrow p\in\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p)$ *and* $<\Gamma,f>\vdash a\not\rightarrow p$*;*
3. *If* $\Gamma\vdash a\rightarrow\sigma\rightarrow q$ *and* $q\rightarrow p\in\Gamma$ *then* $a\rightarrow\sigma\rightarrow q\rightarrow p\in\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p)$ *and* $<\Gamma,f>\vdash a\rightarrow\sigma\rightarrow q\rightarrow p$ *in case* $a\not\rightarrow p\notin\Gamma$ *and either* $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p)=\emptyset$ *or* $f(a,p)=\mathrm{T}$*;*
4. *If* $\Gamma\vdash a\rightarrow\sigma\not\rightarrow q$ *and* $q\not\rightarrow p\in\Gamma$ *then* $a\rightarrow\sigma\rightarrow q\not\rightarrow p\in\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p)$ *and* $<\Gamma,f>\vdash a\rightarrow\sigma\rightarrow q\not\rightarrow p$ *in case* $a\rightarrow p\notin\Gamma$ *and either* $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p)=\emptyset$ *or* $f(a,p)=\mathrm{F}$.

We will also use '$\vdash$' for the relation of support between $<\Gamma,f>$ and propositions. This will cause no confusion, since we will write, e.g., '$<\Gamma, f> \vdash a \rightarrow p$' when '$\vdash$' represents the permission relation where paths are involved, and '$<\Gamma, f> \vdash pa$' when '$\vdash$' represents the support relation. Thus '$\Gamma \vdash a \rightarrow p$' means that $a \rightarrow p \in \Gamma$; '$<\Gamma, f> \vdash pa$' means that $<\Gamma, f>$ yields pa as a conclusion, whether or not it is an *immediate* conclusion, i.e., whether or not $a \rightarrow p \in \Gamma$.

Definition 1 is motivated by ideas from [10]: immediate reasons (reasons contained directly in Γ) are never suspended, and complex reasons are suspended only in case a conflict arises between pros and cons. Inheritance is bottom-up, so that a conflict must consist of opposed reasons whose initial subpaths are permitted up to their final links; such a conflict could be said to consist of *prima facie* pros and cons. To provide for credulity, we allow f to resolve such conflicts in favor of a definite conclusion, positive or negative.

A resolution function f for Γ is *adequate for* Γ if $f(a,p)$ is defined whenever $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset$, and $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p) \neq \emptyset$, and $A \rightarrow p \notin \Gamma$ and $A \not\rightarrow p \notin \Gamma$. (A question (a,p) meeting these conditions could be called an *active conflict.*)

The inheritance that is characterized by Def. 1 is very crude; it is "capricious," and does not allow for preemption. This is illustrated by the following networks.

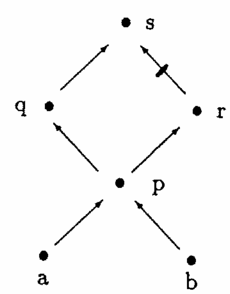

Γ_3

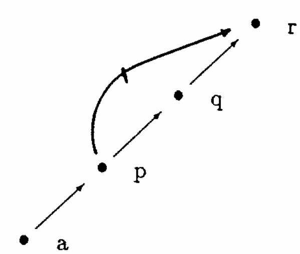

Γ_4

Let $f_3(a,s) = \mathrm{T}$ and $f_3(b,s) = \mathrm{F}$. Let $f_4(a,r) = \mathrm{T}$. Then f_3 is adequate for Γ_3, and $<\Gamma_3, f_3>$ supports sa (and doesn't support $\neg sa$) and also $<\Gamma_3, f_3>$ supports $\neg sb$ (and doesn't support sb). Thus, $<\Gamma_3, f_3>$ is capricious in its failure to treat similar cases differently. Also f_4 is adequate for Γ_4, and $<\Gamma_4, f_4>$ supports ra (and doesn't support $\neg ra$), though its reason for $\neg ra$ is more specific.

The following lemmas bring out some elementary consequences of the definition. It is supposed in all of them that f is adequate for Γ.

Lemma 1 *If $<\Gamma, f> \vdash pa$ and $<\Gamma, f> \vdash \neg pa$ then $a \rightarrow p \in \Gamma$ and $a \not\rightarrow p \in \Gamma$.*

Proof. Cases 3 and 4 of Def. 1 can't apply under the assumptions of the lemma.

Lemma 2 *If $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p) = \emptyset$, then $<\Gamma, f> \nvdash \neg pa$ and $<\Gamma, f> \vdash pa$ iff $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset$. If $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) = \emptyset$ then $<\Gamma, f> \nvdash pa$ and $<\Gamma, f> \vdash \neg pa$ iff $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p) \neq \emptyset$. And if $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset \neq \mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,p)$ then (i) if $pa \in \Gamma$ then $<\Gamma, f> \vdash pa$; (ii) if $\neg pa \in \Gamma$ then $<\Gamma, f> \vdash \neg pa$; (iii) otherwise $<\Gamma, f> \vdash pa$ iff $f(a,p) = \mathrm{T}$ and $<\Gamma, f> \vdash \neg pa$ iff $f(a,p) = F$.*

Proof. Straightforward, by induction on $deg_\Gamma(a,p)$.

Lemma 3 *If* $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset$ *and* $p \rightarrow q \in \Gamma$ *then* $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,q) \neq \emptyset$. *If* $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset$ *and* $p \not\rightarrow q \in \Gamma$ *then* $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,q) \neq \emptyset$.

Proof. Immediate, from Def. 1.

Lemma 4 *If* $<\Gamma, f> \vdash pa$ *then (i) if* $p \rightarrow q \in \Gamma$ *and* $<\Gamma, f> \not\vdash qa$ *then* $<\Gamma, f> \vdash \neg qa$*; and* $\neg qa \in \Gamma$ *or there exists* r *such that* $r \not\rightarrow q \in \Gamma$ *and* $<\Gamma, f> \vdash ra$*; and (ii) if* $p \not\rightarrow q \in \Gamma$ *and* $<\Gamma, f> \not\vdash \neg qa$ *then* $<\Gamma, f> \vdash qa$*; and* $qa \in \Gamma$ *or there exists* r *such that* $r \rightarrow q \in \Gamma$ *and* $<\Gamma, f> \vdash ra$.

Proof. Suppose $<\Gamma, f> \vdash pa$. By Lemma 2, $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,p) \neq \emptyset$. By Lemma 3, if $p \rightarrow q \in \Gamma$ then $\mathrm{R}^{\mathrm{T}}_{\Gamma,f}(a,q) \neq \emptyset$. Then by Lemma 2, if $<\Gamma, f> \not\vdash qa$ then $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,q) \neq \emptyset$ and $f(a,q) \neq T$. Since $\mathrm{R}^{\mathrm{F}}_{\Gamma,f}(a,q) \neq \emptyset$, $\neg qa \in \Gamma$ or there exists r such that $r \not\rightarrow q \in \Gamma$ and $<\Gamma, f> \vdash ra$. Since f is adequate for Γ, $f(a,p)$ is defined; so $f(a,p) = F$, and by Lemma 2, $<\Gamma, f> \vdash \neg qa$. The other half of the lemma is proved in the same fashion.

Lemma 4 characterizes a way in which this type of inheritance is credulous; a reason for a conclusion cannot be overriden without forcing the opposite conclusion.

5.3. A network completeness theorem

We will now show that the model-theoretic notion of consequence that we have defined is equivalent, on the restricted language of networks, to the inheritance definition we gave at the outset. That is, we show that for all acyclic networks Γ and literals A, $\Gamma \mathrel{|\!\!\!\sim} A$ iff $<\Gamma, f> \vdash A$ for all adequate resolution functions f on Γ. The work of the proof is carried out in the following two lemmas.

Lemma 5 *Let* Γ *be acyclic, and* f *be an adequate resolution function on* Γ. *Let the individual constants of* Γ *constitute the domain of* M, *and let* $\mathrm{M} \models_{\mathrm{T}} pa$ *iff* $<\Gamma, f> \vdash pa$, *and* $\mathrm{M} \models_{F} pa$ *iff* $<\Gamma, f> \dashv pa$. *Then* M *is a normal model of* Γ.

Proof. We need to check that M meets the three conditions in the definition of normal models.

1. *(Consistency.)* Suppose $\mathrm{M} \models_{\mathrm{T}} pa \wedge \neg pa$. Then $<\Gamma, f> \vdash pa$ and $<\Gamma, f> \vdash \neg pa$, so by Lemma 1 we have $pa \in \Gamma$ and $\neg pa \in \Gamma$. So M is minimal as regards consistency.

2. *(Anomaly.)* Suppose M' is a model of Γ that also minimizes inconsistency but that $ano_\Gamma(\mathrm{M}') \subset ano_\Gamma(\mathrm{M})$. Then there are some anomalies $<p \rightarrow q, a>$ or $<p \not\rightarrow q, a>$ in $ano_\Gamma(\mathrm{M})$that are absent from $ano_\Gamma(\mathrm{M}'$.) We will derive an inconsistency from this assumption, thereby establishing that M is minimal as regards anomaly.

To begin with, we show by induction on degree that $\mathrm{M} \preceq \mathrm{M}'$; i.e., that if $\mathrm{M} \models_{\mathrm{T}} pa$ then $\mathrm{M}' \models_{\mathrm{T}} pa$ and that if $\mathrm{M} \models_{\mathrm{F}} pa$ then $\mathrm{M}' \models_{\mathrm{F}} pa$. Suppose this holds for questions with degree less than (pa). We can assume without loss of generality that $a \rightarrow p \notin \Gamma$ and $a \not\rightarrow p \notin \Gamma$. Suppose, then, that $<\Gamma, f> \vdash pa$ because $<\Gamma, f> \vdash a \rightarrow \sigma \rightarrow q \rightarrow p$. By inductive hypothesis, $\mathrm{M}' \models_{\mathrm{T}} qa$.

Since $q \to p \in \Gamma$, we must have $\mathrm{M}' \models_{\mathrm{T}} pa$, for otherwise $<a, q \to p>$ would be in $ano_\Gamma(\mathrm{M}')$ but not in $ano_\Gamma(\mathrm{M})$. Now, suppose that $<\Gamma, f> \mid\!\sim pa$ because $<\Gamma, f> \mid\!\sim a \to \sigma \to q \not\to p$. By inductive hypothesis, again $\mathrm{M}' \models_{\mathrm{T}} qa$. Since $q \not\to p \in \Gamma$, we must have $\mathrm{M}' \models \neg pa$, for otherwise $<a, q \not\to p>$ would be in $ano_\Gamma(\mathrm{M}')$ but not in $ano_\Gamma(\mathrm{M})$. This completes the induction.

Assume first that $<p \to q, a>$ is in $ano_\Gamma(\mathrm{M})$ but not in $ano_\Gamma(\mathrm{M}')$. Then we have $\mathrm{M} \models_{\mathrm{T}} pa$, $\mathrm{M} \not\models_{\mathrm{T}} qa$, but $p \to q \in \Gamma$. Since $\mathrm{M} \preceq \mathrm{M}'$, we have $\mathrm{M}' \models_{\mathrm{T}} pa$. Therefore $\mathrm{M}' \models_{\mathrm{T}} qa$. But by Lemma 4, (because $\neg qa \notin \Gamma$) there exists r such that $\mathrm{M} \models_{\mathrm{T}} ra$ and $r \not\to q \in \Gamma$. But then $<r \not\to q, a>$ would be in $ano_\Gamma(\mathrm{M}')$ but not in $ano_\Gamma(\mathrm{M})$, which is impossible. The case in which the anomaly is negative—$<p \not\to q, a>$ is in $ano_\Gamma(\mathrm{M})$ but not in $ano_\Gamma(\mathrm{M}')$—is just the same. This establishes that M is minimal as regards anomaly.

3. *(Information.)* Now suppose M' is a model of Γ that also minimizes inconsistency and anomaly but that $inf(\mathrm{M}') \subset inf(\mathrm{M})$; i.e., $\mathrm{M}' \prec \mathrm{M}$. By the argument with which we began Case 2, $\mathrm{M} \preceq \mathrm{M}'$.

This completes the proof of Lemma 5.

Lemma 6 *Let* M *be a normal model of* Γ, *and construct a resolution function* f *on* Γ *as follows. We let* $f(a,p) = \mathrm{T}$ *if* $\mathrm{M} \models_{\mathrm{T}} pa$ *and* $f(a,p) = \mathrm{F}$ *if* $\mathrm{M} \not\models_{\mathrm{T}} pa$ *and* $\mathrm{M} \models_{\mathrm{F}} pa$; $f(a,p)$ *is undefined otherwise. Then* f *is an adequate resolution function for* Γ, *and for all literals* A, $\mathrm{M} \models_{\mathrm{T}} A$ *iff* $<\Gamma, f> \mid\!\sim A$.

Proof. We prove by induction on $deg_\Gamma(a,p)$ that $\mathrm{M} \models_{\mathrm{T}} pa$ iff there is a path σ enabling pa such that $<\Gamma, f> \mid\!\sim \sigma$, and $\mathrm{M} \models_{\mathrm{F}} pa$ iff there is a path σ enabling $\neg pa$ such that $<\Gamma, f> \mid\!\sim \sigma$, and $\mathrm{M} \models_{\mathrm{T}} pa$ iff $f(a,p) = \mathrm{T}$, and $\mathrm{M} \models_{\mathrm{F}} pa$ iff $f(a,p) = \mathrm{F}$, unless $pa \in \Gamma$.

The basis step is immediate. In the inductive case for a positive formula pa, suppose first that $\mathrm{M} \models_{\mathrm{T}} pa$. If $pa \in \Gamma$ then $<f, \Gamma> \mid\!\sim a \to p$. Suppose $pa \notin \Gamma$. Then there must be a q such that $\mathrm{M} \models_{\mathrm{T}} qa$ and $q \to p \in \Gamma$. Otherwise by removing qa from M we would obtain a model M' such that $\mathrm{M}' \prec \mathrm{M}$ but $ano_\Gamma(\mathrm{M}') \subseteq ano_\Gamma(\mathrm{M})$ and $inc_\Gamma(\mathrm{M}') \subseteq inc_\Gamma(\mathrm{M})$. By the inductive hypothesis, there is a σ enabling qa such that $<\Gamma, f> \mid\!\sim \sigma$. We know $\neg pa \notin \Gamma$ since $\mathrm{M} \models_{\mathrm{T}} pa$ and $pa \notin \Gamma$; therefore $<\Gamma, f> \mid\!\sim \sigma \to p$ iff $f(a,p) = \mathrm{T}$. But $f(a,p) = \mathrm{T}$ because $\mathrm{M} \models_{\mathrm{T}} pa$. Suppose second that $<\Gamma, f> \mid\!\sim \sigma$ for some σ enabling pa. If $pa \in \Gamma$ then $\mathrm{M} \models_{\mathrm{T}} pa$. Otherwise, $\sigma = a \to \tau \to q \to p$. By the inductive hypothesis, $\mathrm{M} \models_{\mathrm{T}} qa$. Since $pa \in \Gamma$, we must also have $f(a,p) = \mathrm{T}$. But then $\mathrm{M} \models_{\mathrm{T}} pa$. The inductive case for a negative formula $\neg pa$ is similar. This completes the proof.

Theorem 1 $\Gamma \parallel\!\sim A$ *iff for all adequate resolution functions* f *on* Γ, $<\Gamma, f> \mid\!\sim A$. *(Here,* Γ *is any acyclic set of formulas of the form* pa, $p \to q$, *or* $p \not\to q$ *and* A *is any literal.)*

Proof. Immediate, from Lemmas 5 and 6.

5.4. The point of the exercise

We feel that the above result provides an illustration, in a simple case, of the sort of theorem that is required in order to properly relate logic and inheritance. The interpretation is faithful and it meets the condition of modularity described below.

But we do not claim that it will be easy to extend this result to capture systems of inheritance that are not capricious, or that are based on preemption, or that are skeptical about conflicting reasons. To construct a model that can account for such things requires further adjustments in the underlying framework of the model theory. We are experimenting with models that incorporate reasons for the conclusions they endorse. These reasons will satisfy general constraints may vary from one model to another; we hope to model different approaches to inheritance with this sort of variation. Our research along these lines, however, is still very tentative.

6. Obstacles to progress

It is often thought that the semantics of inheritance hierarchies can be provided by translating these hierarchies into ordinary nonmonotonic logics. This has been suggested by several writers, including Etherington, Haugh, Lifschitz, and recently Morris.[7] (There is also ongoing work by Gelfond and Pryzmusinska; see [8].)

We feel that the problem of providing semantics for inheritance is unexpectedly complex, and that these attempts have not been fully adequate. We will try to explain this attitude by focusing on Etherington's analysis of credulous inheritance, which is the best-known and most systematic of the developed translations.

6.1. An example: inheritance and default logic

To explain the major source of complexity in Etherington's project of translating inheritance networks into default logic, first consider a straightforward, uncomplicated translation procedure.

Where Γ is a net, first let U_Γ (the universe of Γ) be the set of individual constants occurring in Γ, then let Γ^R be the default theory $\langle W, D\rangle$, where

$$W = \{Pa : a \rightarrow p \in \Gamma\} \cup \{\neg Pa : a \not\rightarrow p \in \Gamma\}$$

$$D = \{(Pa : Qa/Qa) : p \rightarrow q \in \Gamma \;\&\; a \in U_\Gamma\} \cup \{(Pa : \neg Qa/ \neq Qa) : p \not\rightarrow q \in \Gamma \;\&\; a \in U_\Gamma\}$$

This simple translation works well in cases like the net Γ_5 pictured below. (Γ_5 is the notorious Nixon diamond.)

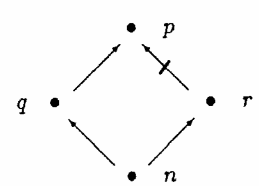

Γ_5

But what does it mean for such a translation to work well *in general*? There must be a natural mapping from the network to the logic, that faithfully represents the inferences sanctioned by the

[7] See [4], [5], [6], [9], [15], [16], [18].

network. More precisely, given a net Γ, let $\overline{\Gamma}$ represent all the *potential conclusions* in Γ—all the statements enabled by any path in Γ. Then we say that the translation is adequate if

- For each default logic extension E of Γ^R, there exists a net Γ' which is an extension of Γ according to inheritance theory, and for which $E \cap [\overline{\Gamma}]^R = [\Gamma']^R$.
- For each extension Γ' of Γ according to inheritance theory, there is a set E which is an extension of Γ^R according to default logic, and for which $E \cap [\overline{\Gamma}]^R = [\Gamma']^R$.

Evidently, the first of these clauses says that the translation is sound with respect to inheritance theory; the second says that it is complete. The restriction in each that $E \cap [\Gamma]^R = [\Gamma']^R$ gives us the appropriate idea of correspondence between default extensions and inheritance extensions.

We will discuss the Nixon diamond Γ_5 as an example. Here $U_{\Gamma_5} = \{n\}$, so Γ_5^R is the default theory

$$\{\{Qn, Rn\}, \{(Qn : Pn/Pn), (Rn : \neg Pn/\neg Pn)\}\}$$

According to default logic, this theory has two extensions: $E_1 = Th(Qn, Rn, Pn)$ and $E_2 = Th(Qn, Rn, \neg Pn)$. In addition, the net Γ_5 itself has two extensions according to inheritance theory, $\Gamma_{5.1} = \{n \rightarrow q, n \rightarrow r, n \rightarrow p\}$ and $\Gamma_{5.2} = \{n \rightarrow q, n \rightarrow r, n \not\rightarrow p\}$. In the case of this net, $\overline{\Gamma_5} = \{n \rightarrow q, n \rightarrow r, n \rightarrow p, n \not\rightarrow p\}$. So each of the two extensions of the default translation do correspond appropriately to an extension of the original inheritance net—and vice versa. For example, $E_1 \cap [\overline{\Gamma_5}]^R = [\Gamma_{5.1}]^R$.

So the simple translation described above indeed works well for the Nixon Diamond. It also works (of course), for any net containing no conflicts at all, and (interestingly) for nets containing only diamond-like conflicts, no matter how deeply nested. It fails, however, for nets in which preemption comes into play.

Consider $\Gamma_6 = \{a \rightarrow p, p \rightarrow q, p \not\rightarrow s, q \rightarrow r, r \rightarrow s\}$.

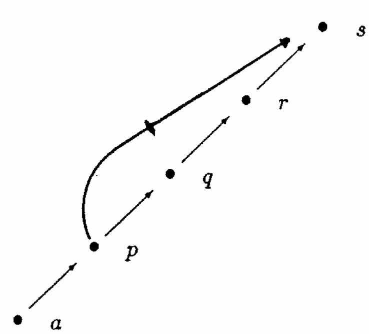

Γ_6

Here Γ_6^R is

$$\{\{Pa\}, \{(Pa : Qa/Qa), (Pa : \neg Sa/\neg Sa), (Qa : Ra/Ra), (Ra : Sa/Sa)\}\}.$$

This default theory has two extensions. But the net Γ_6 itself has only one. One of the two default extensions does not correspond (as above) to the inheritance extensions. So the translation is not sound with respect to inheritance theory. This is the problem that Etherington's more complex solution is designed to solve.

In both inheritance theory and default logic, under the standard translation defined above, we can derive the statements Pa and Ra. At this point, however, we have a choice in default logic, which we does not arise in inheritance theory. We can either use Pa together with the rule $(Pa : \neg Sa/\neg Sa)$ to get $\neg Sa$, or we can use Ra together with the rule $(Ra : Sa/Sa)$ to get Sa. Intuitively, we want to choose the first option, and inheritance theory forces this choice in its treatment of preemption. What forces us, in inheritance theory, to chain off P rather than Q is the interaction of two features of the definition. (1) In inheritance theory, we are able to derive links representing new defaults, as in this case, we can derive the link $p \to r$. (2) Given a conflict between the results of applying two equally applicable original defaults—in this case the P-default and the R-default—inheritance theory uses the *derived* defaults to decide between them.

Now, none of this can be done in default logic, because the logic does not allow us to derive new defaults. Therefore, Etherington has to find some other way of forcing the selection of the P-rule over the R-rule. His solution is to explicitly alter the R-rule so that it cannot apply in cases in which the P-rule can also apply. The new rule is

$$(Ra : [Sa \wedge \neg Pa]/Sa),$$

and of course, when the original translation of the R-rule is replaced by the new one, the resulting theory does have only one extension, which does correspond faithfully to the extension of the network. (The new theory will be semi-normal, no longer normal; but Etherington's results about ordered default theories guarantee that it will have an extension.)

However, this solution is not ideal, for at least the following reasons.

(1) The mapping of nets into semi-normal default theories has only been presented by examples; we have not yet seen a *general* translation scheme. We need a procedure for translating any net Γ into a semi-normal theory Γ^E. And for each net, the translation must be adequate, in the sense we have described.

A major obstacle in working out such a mapping is that *subsumption is not extension invariant.* Whether one statement represents more specific information than another in a given extension, depends on what additional statements are present in that extension. A statement can represent more specific information than another in one extension, but not in every extension.

To see this, consider $\Gamma_7 = \{a \to p,\ p \not\to r,\ a \to q\ q \to r,\ p \to t\ t \to u\ u \to q,\ p \to s\ s \not\to u\}$.

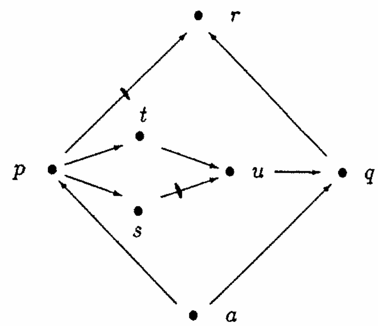

Γ_7

In inheritance theory, the different specificity relations among statements in different extensions can be computed along with the extensions themselves, varying from one extension to another. Thus, inheritance theory naturally gives us three extensions in this net. (See [24] for details.)

However, if one tries to map a network into a seminormal default theory, the different specificity relations among statements in different extensions have to be encoded into the statement of the default rules from the start. Ask youself: "Is the P-rule in Γ_7 more specific than the Q-rule or not?" That is, should the Q-rule be translated simply as

$$(Qa : Ra/Ra),$$

which would put P on a par with Q, or should it be translated as

$$(Qa : [Ra \wedge \neg Pa]/Ra)$$

which would give P priority? The answer is that neither translation is correct! According to the view of preemption adopted in inheritance theory, P turns out to be more specific than Q just in case one choses $a \rightarrow p \rightarrow t \rightarrow u$ over the competing $a \rightarrow p \rightarrow s \not\rightarrow u$. In order to translate this net adequately into default logic, therefore, we would have to map the Q-rule $(q \rightarrow r)$ into something like

$$(Qa : [Ra \wedge (Ua \supset \neg Pa)]/Ra),$$

which gives the P-rule priority in exactly those cases in which Ua holds—which is what is desired in this particular net. Notice how, even in this simple case, it becomes complicated to encode the different priority relations among different rules in different extensions into the initial statement of those rules. Moreover, it is virtually impossible to derive the encoding without knowing what the relevant set of extensions is going to be. And of course, if one already needs to precalculate the set of extensions associated with a network in order to code it properly into its a default theory, then there is only a very weak sense in which the resulting default theory can be said to "give the meaning" of the network.

(2) The translation is not modular. We want a uniform translation for each link, a translation that is independent of the context in which the link occurs. Otherwise, updating a knowledge base will become a very complicated operation. If the translation of any particular link cannot be determined independently of the context in which it occurs, then when one updates a net by adding a new link, it may be incorrect to update the translation by simply adding the translation of this new link. The new link will alter the context in which the original links occur; and so the translations of these links may have to be adjusted to reflect the change in context.

As an example, consider the translation given above of Γ_5 (the Nixon Diamond). Suppose we update this net by adding to it the link $q \rightarrow r$. (Of course, it destroys the plausibility of the original interpretation to suppose that Quakers are Republicans, but ignore that for now.) The resulting net no longer has two extensions. But how should we update the default theory that results from translating this net? It is not enough simply to add to the set of default rules the natural translation of this new link—$(Qn : Rn/Rn)$. If we did only that, the resulting default theory would still have two extensions; so it would no longer be sound with respect to inheritance theory. Instead, we have to realize that, in the new context, the original translation of $r \rightarrow p$ is no longer adequate; because of the addition of the new link, the translation of this original link must now be changed to $(Rn : [\neg Pn \wedge \neg Qn]/\neg Pn)$.

As you can see, the update operation defined on the net was quite simple; it was just a matter of adding a new link. However, because the translation is non-modular, the corresponding update operation in the default theory is more complicated; in addition to adding the translation of the new link, it is necessary to adjust the translation of links already present to fit the new context. And, of course, in general the adjustments can become much more complex.

Here, we have given space only to Etherington's approach, and have summarized the difficulties. The points we make here are treated in much more detail, with more examples, and are generalized to a variety of translation proposals, in [12].

6.2. Where we stand

What we've said so far applies only to the attempts to translate credulous inheritance theories. And this is the easiest translation task, because of the credulous bias of most current nonmonotonic logics (and of default logic in particular). Here, as noted, the logic works well in the case of diamond-like conflicts—serious difficulties do not arise until one tries to incorporate preemption. In the case of skeptical inheritance theories, the logics do not even agree in the case of diamond-like conflicts; the only case in which there is a coincidence is with networks that are entirely conflict free.

The problems we have found in Etherington's program were meant to illustrate general difficulties facing a number of different analyses of inheritance in terms of standard nonmonotonic logics. For example, the analysis in [18] (which also uses default logic) is also non-modular—although here, the problem does not arise in the (nonnormal) default rules that Morris uses, but in the additional axioms he has to add in order to make his default rules have the right effect. Proposals relying on prioritized circumscription also run into similar problems: priorities among predicates have to be fixed in advance; so there is no way of allowing the priorities to vary from minimal model to minimal model.[8]

In any event, we expect the next few years of research to produce formally adequate solutions to the translation problem. Our attempt in the present paper to build more flexible models based on the four-valued interpretations that characterize nonmonotonic inheritance is an idea that in some ways is like Lifshitz' proposal, and that certainly is motivated by similar considerations. And a number of new formalisms are emerging for nonmonotonic reasoning which attempt to mix the expressive power of full default logic with the implicit preference for more specific arguments found in inheritance theory. Among these are the systems of Delgrande, Geffner and Pearl, Horty, Loui, and Poole.[9]

Horty has recently proved that there is an adequate and modular translation of certain inheritance theories into the logic he describes; and also, that this logic is, in a certain sense, sound with respect to default logic. As far as we know, no exact correspondence has yet been found between inheritance theory and the systems of Delgrande, Geffner and Pearl, Loui, and Poole. But things are moving very quickly in this area of research, and we should soon be in a position to debate the merits of several formally correct solutions.

But the problems we have raised here are real, and can only be overcome by modifications in the logics that run deep. In general, the trend seems to be one that was foreseen long ago by Jon Doyle;[10] to replace the descriptive, truth-oriented bias of familiar logical theories with more normative, value-oriented conceptions, such as the idea of priorities on reasons for conclusions. Speaking now as logicians, we welcome this development. Of course, logic has its own internal research momentum, but like any other living field of inquiry it can always use new ideas. Since the theories that are emerging from the theoretical A.I. community are rigorously formulated and computationally significant, they may well prove to be valuable and interesting additions to the world of logic.

7. Acknowledgment

This material is based on work supported by the National Science Foundation under Grant No. IST-8700705 (Logical Foundations for Inheritance and Knowledge Update).

[8] In [16] there is a new formulation of prioritized circumscription, which allows one to reason about priorities within the theory. This is the kind of apparatus that is needed to solve the translation problem; we are not yet familiar enough with the details of Lifshitz' proposal to see if in fact it allows a solution to the problem.

[9] See [2], [7], [13], [17], [19].

[10] See [3] for a general presentation.

Bibliography

[1] N. Belnap. "A useful four-valued logic." In J. Dunn and G. Epstein (eds.), *Modern uses of multiple-valued logic*, D. Reidel, 1977, pp. 8–37.

[2] J. Delgrande, "An approach to default reasoning based on a first-order conditional logic: revised report." *Artificial Intelligence*, vol. 36 (1988), pp. 63–90.

[3] J. Doyle. Some theories of Reasoned Assumptions. Department of Computer Science, Carnegie Mellon University, 1983. Technical Report No. CMU-CS-83-125.

[4] D. Etherington and R. Reiter, "On inheritance hierarchies with exceptions." In *Proceedings of AAAI-83*, Morgan Kaufmann (1983), pp. 104–108.

[5] D. Etherington, "Formalizing nonmonotonic reasoning systems." *Artificial Intelligence*, vol. 31 (1987), pp. 41–85.

[6] D. Etherington, *Reasoning with Incomplete Information.* Morgan Kaufmann (1988), viii+240 pp.

[7] H. Geffner and J. Pearl, "A framework for reasoning with defaults." Technical Report R-94-III, Computer Science Department, University of California at Los Angeles (1988), 17 pp.

[8] M. Gelfond and H. Przymusinska, "Inheritance hierarchies and autoepistemic logic." Forthcoming Technical Report, Computer Science Department, University of Texas at El Paso.

[9] B. Haugh, "Tractable theories of multiple defeasible inheritance in ordinary nonmonotonic logics." In *Proceedings of AAAI-88*, Morgan Kaufmann (1988), pp. 421–426.

[10] J. Horty, R. Thomason, and D. Touretzky, "A skeptical theory of inheritance in nonmonotonic semantic nets." Department of Computer Science, Carnegie Mellon University, 1987. Technical Report No. CMU-CS-87-175. Forthcoming in *Artificial Intelligence.*

[11] J. Horty, R. Thomason, and D. Touretzky, "A skeptical theory of inheritance in nonmonotonic semantic nets." In *AAAI-87* (Proceedings of the Sixth National Conference on Artificial Intelligence). Volume 2, Morgan Kaufmann, Los Altos, 1987, pp. 358–363.

[12] J. Horty, "Nonmonotonic inheritance and nonmonotonic logic." Forthcoming Technical Report, Computer Science Department, Carnegie Mellon University.

[13] J. Horty, "Credible extensions of defeasible theories." Forthcoming Technical Report, Computer Science Department, Carnegie Mellon University.

[14] K. Konologe, "Hierarchic autoepistemic theories for nonmonotonic reasoning." In *Proceedings of AAAI-88*, Morgan Kaufmann (1988), pp. 439–443.

[15] V. Lifschit
workshop, N

[16] V. Lifshitz
tion." *Journ*

[17] R. Loui, "I
ligence, vol.

[18] P. Morris,
vol. 35 (198

[19] D. Poole,
Proceedings

[20] R. Thomaso
nets." Depa
No. CMU-C

[21] R. Thomaso
nets." In Z.
of the Secon
Holland, An

[22] R. Thomaso
Forthcoming

[23] D. Touretzk
84, Morgan

[24] D. Touretzk
1986.

[25] D. Touretzk
monotonic n
Los Altos, 1

[15] V. Lifschitz, "Some results on circumscription." *Proceedings of the non-monotonic reasoning workshop*, New Paltz, NY. AAAI (1984), pp. 151–164.

[16] V. Lifshitz, "Circumscriptive theories: a logic-based framework for knowledge representation." *Journal of Philosophical Logic,* vol. 17 (1989), forthcoming.

[17] R. Loui, "Defeat among arguments: a system of defeasible inference." *Computational Intelligence*, vol. 3 (1987), pp. 100–106.

[18] P. Morris, "The anomalous extension problem in default reasoning." *Artificial Intelligence*, vol. 35 (1988), pp. 383–399.

[19] D. Poole, "On the comparison of theories: preferring the most specific explanation." In *Proceedings of IJCAI-85*, Morgan Kaufmann (1985), pp. 144–147.

[20] R. Thomason, J. Horty, and D. Touretzky, "A calculus for inheritance in monotonic semantic nets." Department of Computer Science, Carnegie Mellon University, 1986. Technical Report No. CMU-CS-86-138.

[21] R. Thomason, J. Horty, and D. Touretzky, "A calculus for inheritance in monotonic semantic nets." In Z. Ras and M. Zemankova (eds.), *Methodologies for intelligent Systems* (Proceedings of the Second International Symposium on Methodologies for Intelligent Systems), North-Holland, Amsterdam, 1987, pp. 280–287.

[22] R. Thomason and J. Aronis, "Hybridizing nonmonotonic inheritance with theorem proving." Forthcoming.

[23] D. Touretzky, Implicit ordering of defaults in inheritance systems. In *Proceedings of AAAI-84*, Morgan Kaufmann (1984), pp. 322–325.

[24] D. Touretzky, *The mathematics of inheritance systems.* Morgan Kaufmann, Los Altos, CA, 1986.

[25] D. Touretzky, J. Horty, and R. Thomason, "A clash of intuitions: the current state of nonmonotonic multiple inheritance systems." In *Proceedings of IJCAI-87*, Morgan Kaufmann, Los Altos, 1987, pp. 476–482.

Lecture Notes in Artificial Intelligence

Subseries of Lecture Notes in Computer Science

Edited by J. Siekmann

J. Horty

346

M. Reinfrank J. de Kleer M. L. Ginsberg
E. Sandewall (Eds.)

Non-Monotonic Reasoning

2nd International Workshop
Grassau, FRG, June 13–15, 1988
Proceedings

Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo

ɔf of
f titles
e
y vis-
ecture
ture
rial
who
ed-

ted
we
ding
unity.