

A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems

David S. Touretzky
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

John F. Horty
Philosophy Department
University of Maryland
College Park, MD 20742

Richmond H. Thomas on
Linguistics Department
University of Pittsburgh
Pittsburgh, PA 15260

Abstract: Early attempts at combining multiple inheritance with nonmonotonic reasoning were based on straight forward extensions of tree-structured inheritance systems, and were theoretically unsound. In *The Mathematics of Inheritance Systems*, or *TMOIS*, Touretzky described two problems these systems cannot handle: reasoning in the presence of true but redundant assertions, and coping with ambiguity. *TMOIS* provided a definition and analysis of a theoretically sound multiple inheritance system, accompanied by inference algorithms. Other definitions for inheritance have since been proposed that are equally sound and intuitive, but do not always agree with *TMOIS*. At the heart of the controversy is a clash of intuitions about certain fundamental issues such as skepticism versus credulity, the direction in which inheritance paths are extended, and classical versus intuitive notions of consistency. Just as there are alternative logics, there may be no single "best" approach to nonmonotonic multiple inheritance.

1. Introduction

Early attempts at combining multiple inheritance with exceptions were based on straightforward extensions of tree-structured inheritance systems, and were theoretically unsound. Two well-known examples are FRL [Roberts and Goldstein 1977] and NETL [Fahlman 1979]. In *The Mathematics of Inheritance Systems*, or *TMOIS*, Touretzky described two classes of problems that these systems cannot handle [Touretzky 1986]. One involves reasoning with true but redundant assertions; the other involves ambiguity.

TMOIS provided the definition and analysis of a theoretically sound multiple inheritance system, along with some inference algorithms based on parallel marker propagation. Other definitions for inheritance have since been proposed in [Sandewall 1986] and [Horty *et al.* 1987] that are equally sound and intuitive, but do not always agree with the system defined in *TMOIS*. At the heart of the controversy is a clash of intuitions about certain fundamental issues in inheritance reasoning. In this paper we catalog the issues, map out a design space, and describe some interesting properties that result from various choices of definitions. Just as there are alternative logics, there may be no single "best" approach to nonmonotonic multiple inheritance.

This material is based on work supported by the National Science Foundation under Grant No. IST-8516313.

2. Basic Terminology

An **inheritance network** is a labeled directed graph whose nodes represent individuals and classes, and whose links denote various relations between the nodes. The primary relation is the IS-A link, which is written as $x \rightarrow y$. Links of form $x \not\rightarrow y$ are called IS-NOT-A links. If a is an individual and p and q are classes, the links $a \rightarrow p$ and $a \not\rightarrow q$ have straightforward logical interpretations as the sentences $p(a)$ and $\sim q(a)$, respectively. Links between classes, such as $p \rightarrow q$ and $p \not\rightarrow q$, have no fixed logical interpretation; they *might* represent sentences in a first order, default, or nonmonotonic logic, depending on the type of inheritance system in which they appear.

A **tree-structured inheritance system** is one whose nodes and IS-A links form a tree. In contrast, a **multiple inheritance system** does not restrict nodes to having at most one immediate superior; typically the only requirement is that the inheritance graph (or at least the IS-A subgraph) must be acyclic. Acyclic multiple inheritance graphs are sometimes referred to as "tangled hierarchies" [Fahlman 1979].

A **bipolar inheritance system** is one that contains both positive (\rightarrow) and negative ($\not\rightarrow$) links. A **unipolar system** contains only positive links, and therefore cannot directly express negative statements such as "Clyde is not a cabbage." (One can represent negative statements indirectly in a unipolar system by creating a positive link to a "non-cabbage" node, but unipolar inference algorithms won't support it properly.) NETL and *TMOIS* are examples of bipolar systems. FRL and the type systems of object-oriented programming languages are unipolar.

A **nonmonotonic inheritance system** permits exceptions to inherited properties. NETL and *TMOIS* rely on bipolar links to achieve nonmonotonicity: negative links are used to override positive inherited information, and positive links to override negative inherited information. Other systems use explicit exception links for this purpose [Borgida 1985], [Etherington 1987], [Fahlman *et al.* 1981]. Bipolarity and nonmonotonicity are independent properties. The type systems of object-oriented programming languages are unipolar monotonic, but frame systems such as FRL are unipolar non-monotonic. In a recent technical report, we describe a bipolar monotonic system [Thomason *et al.* 1986].

A **homogeneous inheritance system** is either monotonic (all links are strict) or thoroughly nonmonotonic (all

links are defeasible.) Heterogeneous inheritance systems have also been proposed, e.g., in [Etherington 1987]; these would contain a mixture of strict and defeasible links. In this paper we are concerned with bipolar, nonmonotonic, homogeneous multiple inheritance systems, which we feel are the most interesting homogeneous class. Heterogeneous systems are not yet well understood, and are not considered here.

The theory of an inheritance network is a fixedpoint consisting of paths of form $x_1 \rightarrow \dots \rightarrow x_n$ or $x_1 \rightarrow \dots \rightarrow x_{n-1} \nrightarrow x_n$ which are generated by inference rules. (This definition relies on the fact that, due to the nature of the inference rules, \nrightarrow may appear only as the last link of a path.) Inference rules derive new paths from existing ones, starting with the paths of length 2 that are the links of the net. A fixedpoint of a nonmonotonic network is called an extension rather than a theory, because theories are assumed to be deductive and therefore unique. Nonmonotonic inference is not purely deductive; nonmonotonic nets may have multiple extensions.

The set of conclusions associated with an extension is the set of statements of form $x_1 \rightarrow x_n$ if the extension contains a path of form $x_1 \rightarrow \dots \rightarrow x_n$, or $x_1 \nrightarrow x_n$ if the extension contains a path of form $x_1 \rightarrow \dots \rightarrow x_{n-1} \nrightarrow x_n$. It is possible for two distinct extensions to be associated with the same set of conclusions, if they arrive at those conclusions through different inference paths.

3. Shortest Path Reasoners

In the past, multiple inheritance systems have usually been defined informally, e.g., by specifying a procedure for traversing an inheritance graph in response to a query. These reasoners, which use either breadth-first or depth-first searches, are called shortest path reasoners, because the answer they return is always the one with the shortest inference path from the starting node. A typical example is NETL, which was designed in conjunction with special parallel hardware that could perform breadth-first searches quickly.

Shortest path algorithms are efficient, and they are provably correct for tree-structured inheritance systems with or without exceptions. They are also appropriate for multiple inheritance systems that do not admit exceptions. However, when multiple inheritance is combined with exceptions, the definition of correct inheritance reasoning becomes much more complex. Although it is still unknown whether there is a single "best" reasoning strategy for multiple inheritance with exceptions, the following two problems clearly eliminate shortest path reasoning as a candidate.

3.1. Reasoning With Redundant Statements

Figure 1 illustrates the problem of reasoning in a net containing true but redundant statements, which can only occur in a multiple inheritance system. Clyde is an elephant and elephants are gray, but Clyde is also a royal




Figure 1: A network with a redundant link.

elephant, and royal elephants are not gray. The redundant IS-A link from Clyde to elephant causes problems for simple-minded inference algorithms because in the presence of such links inheritance paths are no longer independent. In particular, shortest path reasoning algorithms no longer apply.

The intuition underlying inheritance with exceptions is that claims about subclasses are more specific, and so can incorporate information about exceptional cases. Therefore, subclasses should be allowed to override superclasses. Clyde is not gray because royal elephants aren't gray; the information from royal elephant overrides the information from elephant despite the presence of a direct redundant link from Clyde to elephant.

3.2. Reasoning in the Presence of Ambiguity

Figure 2 illustrates the problem of reasoning in the presence of ambiguity. This familiar example is known as the Nixon diamond. Quakers are typically pacifists while Republicans typically are not. Nixon is both a Quaker and a Republican. Is he a pacifist or isn't he? One might choose to draw no conclusion, or one might generate two extensions as *TMOIS* does, with Nixon being a pacifist in one and a non-pacifist in the other. The crucial point is that the ambiguity must be recognized.

Simple depth-first search algorithms, which work perfectly well for tree-structured networks, have no way to recognize the ambiguities that can arise in multiple inheritance situations. Shortest path reasoners will either choose one of the two possible conclusions about Nixon arbitrarily, or else report that both conclusions are true at the same time, i.e., that Nixon simultaneously is and is not a pacifist. Neither behavior is theoretically sound. The inability to recognize ambiguity is another reason why shortest path algorithms are unacceptable for reasoning about situations involving nonmonotonic multiple inheritance.




Figure 2: The Nixon diamond is the canonical example of an ambiguous net.

4. A Clash of Intuitions

The two problems described in section 3 were solved in *TMOIS* by introducing a measure called the **inferential distance ordering** as a replacement for path length. The precise definition of inferential distance is highly technical, but it essentially states that an inference path $A \rightarrow \dots \rightarrow B \not\rightarrow P$ should preempt a path $A \rightarrow \dots \rightarrow C \rightarrow P$ iff there is a path $A \rightarrow \dots \rightarrow B \rightarrow \dots \rightarrow C$ establishing for A that B is a subclass of C .

Although inferential distance is important for distinguishing inheritance systems from other nonmonotonic reasoning systems such as default logic [Touretzky 1984], it does not by itself yield a unique definition of nonmonotonic multiple inheritance. There are a number of other choices to be made. Touretzky built into *TMOIS* certain assumptions about reasoning that other investigators, approaching inheritance with a different set of intuitions, may not share. In the remainder of this section we describe four areas in which intuitions have clashed. It should be stressed that all of these systems treat redundancy and ambiguity in the correct way, and appear both theoretically sound and intuitively appealing. Nonetheless, we will describe cases in which they give conflicting results.

4.1. Skeptical vs. Credulous Reasoning

A skeptical reasoner refuses to draw conclusions in ambiguous situations.¹ A skeptical reasoner would therefore offer no opinion as to whether Nixon was or was not a pacifist. Credulous reasoners, on the other hand, try to conclude as much as possible. To avoid inconsistency, a credulous reasoner must generate multiple extensions so that contradictory conclusions can be isolated from one another. *TMOIS* describes such a reasoner. In one of the two credulous extensions of figure 2 the inference path $NIXON \rightarrow QUAKER \rightarrow PACIFIST$ is admissible; in the other, the path $NIXON \rightarrow REPUBLICAN \rightarrow PACIFIST$ is admissible.

¹This was the attitude toward reasoning of the first philosophical skeptics; they refused to draw a conclusion in cases in which any reason could be produced for the opposite conclusion.




Figure 3: Cascaded ambiguities.

Since skeptical reasoning always generates a unique extension, there is no need to consider sets of extensions. We expect skeptical reasoning algorithms to be simpler and more efficient than credulous ones; in our investigations so far, this has indeed been the case. On the other hand, one advantage of credulity is that after generating all permissible extensions they can be examined for interesting properties.

Skeptical inheritance as defined in [Horty *et al.* 1987] is not equivalent to taking the intersection of all credulous extensions. One reason is that, in the presence of multiple ambiguities, skeptical inference may permit certain paths that are not permitted in all credulous extensions. Consider figure 3. Nixon's pacifism is ambiguous here, and consequently, so are his feelings about the military. But a simple skeptical reasoner would not recognize the second ambiguity. The reason is that since it is ambiguous whether or not Nixon is a pacifist, there are no valid paths from Nixon to pacifist to anti-military in the skeptical extension. Therefore the path from Nixon to football fan to not-anti-military is unopposed, so one may conclude that Nixon is not anti-military. This is the approach recommended in [Horty *et al.* 1987].

It might be preferable for ambiguity to be propagated from pacifist to anti-military, based on the following argument: we are reserving judgement about whether Nixon is a pacifist or not, but conceivably he could be one. In that case he could conceivably be anti-military. Being a football fan he could also conceivably be not anti-military, so we should reserve judgement on whether he is anti-military.

An ambiguity-propagating version of skepticism would be more compatible with credulous reasoning, in that the skeptic never draws conclusions about matters that would be decided differently in different credulous extensions. The goal of ambiguity-propagating skepticism is to produce a

single extension whose conclusions are consistent with the intersection of all credulous extensions. It is not yet known whether this can be done without computing all the credulous extensions.

4.2. Upward vs. Downward View of Reasoning

The most common intuitive model of inheritance has properties flowing downward from classes to their subclasses and instances—though the flow can be interrupted by an exception. This downward view is the one promoted in *TMOIS* and, informally, in *FRL* and *NETL*. An opposing view of inheritance has inference working upwards. This conception, which is explicitly chosen in [Horty *et al.* 1987], stresses arguments (or proof sequences), and imagines them being constructed from the bottom up.³ The upward and downward views aren't always compatible, as the following two examples illustrate.

4.2.1. Coupling in Downward Reasoners

Figure 4 demonstrates an important difference between upward and downward credulous reasoners. (In this section we consider only credulous reasoners; the skeptical approach would generate a single extension with no interesting conclusions about A or B.) A downward credulous reasoner produces two extensions from figure 4, due to the ambiguity about whether B's are E's. One extension contains $B \rightarrow C \rightarrow E$ and $A \rightarrow B \rightarrow C \rightarrow E$ while the other contains $B \rightarrow D \nrightarrow E$ and $A \rightarrow B \rightarrow D \nrightarrow E$. Notice that in each extension, the conclusion about A's being E's is identical to the conclusion about B's being E's. The property that a subclass is always in agreement with its superclasses (in the absence of explicit exceptions) is known as coupling. Downward reasoners necessarily produce coupled theories because the only properties a node can inherit are those of its superiors. Upward reasoners are not so constrained.

In an upward credulous reasoner the arguments made about A in figure 4 need not depend on conclusions that were drawn about B, since properties are not flowing down from B to A. So in addition to the above two extensions, an upward reasoner could produce two uncoupled extensions. **One would contain $B \rightarrow C \rightarrow E$ and $A \rightarrow B \rightarrow D \nrightarrow E$; the other would contain $B \rightarrow D \nrightarrow E$ and $A \rightarrow B \rightarrow C \rightarrow E$.**

Intuitively, coupling appears to be a desirable property for reasoners to have. It may be possible to constrain upward reasoners to produce only coupled theories, but this could compromise their ability to reach unambiguous conclusions where downward reasoners do not, a property known as opportunistic inference.

4.2.2. Opportunism in Upward Reasoners

Node B is ambiguous in figure 5 for the same reason it's ambiguous in figure 4: there are two mutually contradic-

² Note that although both *FRL* and *NETL* present properties as flowing downward from classes to individuals, in fact their search algorithms proceed in the opposite direction, starting with the individual and proceeding up the ISA hierarchy to increasingly general superclasses. Only *TMOIS* generates inference paths from top to




Figure 4: Reasoners that exhibit coupling treat A and B identically.




Figure 5: B is ambiguous about property E; should A be?

tory paths $B \rightarrow C \rightarrow E$ and $B \rightarrow D \nrightarrow E$. For now let us consider only skeptical reasoners. No skeptical reasoner should draw a conclusion about whether B's are E's. But what should it conclude about A's being E's? In a downward skeptical reasoner, since B is ambiguous and therefore neither an E nor a non-E, there is no property for A to inherit. So a downward reasoner would draw no conclusion about whether A's are E's.

In an upward skeptical reasoner, the link $A \nrightarrow D$ rules out the potential inference path $A \rightarrow B \rightarrow D \nrightarrow E$, leaving the path (or argument) $A \rightarrow B \rightarrow C \rightarrow E$ to go through unopposed.

The upward reasoner seizes on any chance to prove something, in this case that A's are E's, even though B remains ambiguous on this point. Because upward reasoners are opportunistic, they sometimes reach conclusions in situations where a downward reasoner would abstain.

Upward and downward *credulous* reasoners also behave differently with respect to figure 5. A credulous *upward* reasoner would generate two extensions: one in which B's were E's, and one in which they were not. A's would be E's in both extensions. *TMOIS*, a credulous *downward* reasoner,

would also generate two extensions, but the extension containing $B \rightarrow D \wedge E$ would contain no conclusion at all about A 's, since the potential path $A \rightarrow B \rightarrow D \wedge E$ is cancelled by the $A \wedge D$ link.

4.3. On-Path vs. Off-Path Preemption

Figure 1, which contains a redundant link, supports two conflicting inference paths: $CLYDE \rightarrow ELEPHANT \rightarrow GRAY-THING$ and $CLYDE \rightarrow ROYAL-ELEPHANT \wedge GRAY-THING$. The latter path is said to preempt⁸ the former. Since royal elephant is on the path from Clyde to elephant, presumably it has more specific information about Clyde than elephant does. The redundant direct link from Clyde to elephant caused problems for early nonmonotonic inheritance reasoners precisely because shortest path algorithms don't correctly implement preemption.

An inheritance reasoner may adopt either of two versions of preemption, which we call *on-path* and *off-path preemption*. On-path preemption, first described in *TMOIS*, is the stricter approach: one path may preempt another only if the preempted path contains a redundant link (or "level skip," as it's called in *NETL*) that would short-circuit part of the preemptor. If a node interrupts a redundant link, as african-elephant interrupts the link from Clyde to elephant in figure 6, the preemption relation no longer holds. Figure 6 generates the two conflicting paths $CLYDE \rightarrow ROYAL-ELEPHANT \wedge GRAY-THING$ and $CLYDE \rightarrow AFRICAN-ELEPHANT \rightarrow ELEPHANT \rightarrow GRAY-THING$. Even though royal elephant is on the path from Clyde to elephant, preemption doesn't hold because royal elephant is not on any path from Clyde through African elephant to elephant; in other words, figure 6 contains no redundant links.

Choosing the downward view of reasoning for a moment, we note that royal elephants are not gray while African elephants are gray. Since neither class is a subclass of the other, and Clyde is an instance of both, he could inherit either one's grayness property. Clyde's grayness is therefore ambiguous, so the network should have two extensions. A similar argument holds if we choose the upward view of inheritance.

Sandewall argues that on-path preemption in *TMOIS* is too restrictive, citing figure 6 as an example where our intuitions suggest that Clyde is unambiguously not gray [Sandewall 1986]. He proposes a more permissive rule, one which we call off-path preemption. The argument for off-path preemption begins with the fact that both royal elephant and African elephant's paths to gray thing must go through elephant. The fact that African elephants are typically gray shouldn't influence Clyde, because African elephants aren't explicitly gray, they merely inherit grayness from elephant. Royal elephant overrides the elephant node's grayness property with more specific information.

This argument does lead to the desired conclusion that Clyde is not gray in figure 6. Off-path preemption in gen-




Figure 6: Sandewall's example.




Figure 7: A possible counterexample to Sandewall's net, with the same topology.

eral is more permissive (more willing to let one path override another rather than cause an ambiguity) than on-path preemption. But is it more intuitive?

Consider figure 7, which has the same topology as figure 6 but different node labels. Marines are typically men, and most men drink beer. Chaplains are typically men but they typically don't drink beer. George is a marine chaplain. In this situation we are less certain that George doesn't drink beer, which the off-path version of preemption would insist is the case. Marines and chaplains are quite different from each other; possibly even antithetical. Both types of individual are rather different from the typical man. It is therefore difficult to predict which properties George will inherit from chaplain and which from marine. Perhaps the specific knowledge that chaplains don't drink beer should take precedence over what we know about men, but perhaps not. The most relevant missing bit of information is the rate of beer drinking among marines. If this is far higher than the rate of abstention among chaplains, one would be better off assuming George is a beer drinker than

⁸Instead of "preemption," the term "preclusion" was USED in *TMOIS*. The concept is the same.

not.⁴

Like Sandewall's system, the skeptical reasoner described in [Horty *et al.* 1987] uses off-path rather than on-path preemption. A permissive view of preemption allows a skeptical system to reach conclusions where a strict on-path approach would see an ambiguity. On the other hand, a credulous reasoner that chose off-path preemption would simply produce a subset of the extensions allowable with on-path preemption; for credulous reasoners it may pay to take a more cautious, strict approach.

4.4. Classical vs. Intuitive Notions of Consistency

According to classical logic, if it is possible to derive any pair of statements p and $\sim p$ then the initial axioms are inconsistent and the only theory that the axioms generate is the one containing every statement. This theory of course destroys any useful information that may be contained in the axioms.

Graph-oriented approaches to reasoning avoid this problem, since a statement A cannot be derived from a set of statements T unless T creates a path supporting A . In fact, even strictly monotonic inheritance networks are non-classical in this respect, as we show in [Thomason *et al.* 1986]. Figure 8a is a classically inconsistent net: it contains the statements $p(a)$ and $\sim p(a)$. Such a net might represent the sentences "Clyde is an elephant" and "Clyde is not an elephant." Although we can't use it to derive all possible statements, the logical inconsistency of the network is easily apparent.

This side-effect of graph-based reasoning is clearly desirable, since it enables graceful degradation of information retrieval in the presence of inconsistent data. Here, classical logic is inadequate as a theory of inheritance reasoning, and should be replaced with a more adequate approach. In [Thomason *et al.* 1966], we show that the simple four-valued logic of [Belnap 1977a], [Belnap 1977b] is adequate for characterising monotonic inheritance.

In other cases, classical logic is overly liberal: figure 8b is not classically inconsistent. This figure could represent the sentences "elephants are gray" and "elephants are not gray." **The two links of the figure are expressed in predicate calculus as $(\forall x)[p(x) \rightarrow q(x)]$ and $(\forall x)[p(x) \rightarrow \sim q(x)]$.** As long as there are no instances of elephants, there is no logical contradiction here, even though the topology of the network is the same as that of figure 8a. *TMOIS* labels figure 8b inconsistent because even if there are no instances of p 's, it makes no sense to simultaneously believe both "the typical p is a q " and "the typical p is not a q " Another argument for the inconsistency of 8b is that it simplifies things if networks with identical topologies are treated identically.

The intuitive notion of consistency put forth in *TMOIS*,

⁴In fact, if marines are unusually avid consumers of beer, one might argue that marine should have a direct link to beer drinker. Such a link would be redundant as far as the marine node itself is concerned, but it would supply additional evidence in favor of beer drinking for particular instances of marine, such as George.



Figure 8: Two inconsistent nets.

motivated by the desire to talk about vaguely-specified concepts such as "the typical p ," is at odds not only with classical logic but with the default and nonmonotonic logic formulations of inheritance as well. For example, figure 8b is usually expressed in default logic [Reiter 1980] as:

$$\frac{p(x) : q(x)}{q(x)} \quad \text{and} \quad \frac{p(x) : \sim q(x)}{\sim q(x)}$$

If multiple instances of p 's exist, some may be inferred (via the first default) to be q 's, and others (via the second default) not to be q 's; in any case no contradiction will arise. The same is true in nonmonotonic logic, where figure 8b would be expressed as:

$$p(x) \wedge M[q(x)] \rightarrow q(x)$$

$$p(x) \wedge M[\sim q(x)] \rightarrow \sim q(x)$$

A conservative reasoner could refrain from assigning any interpretation to a net containing even a single inconsistency. But since an inconsistency in one part of the net doesn't necessarily affect inference in other parts, it may be possible to reason around the inconsistency and in some cases produce useful inferences.

5. Conclusion

The study of inheritance reasoning has rapidly moved through stages similar to ones that deductive logic has passed through, at a much more glacial pace. Reflection on reasoning practices has produced a large body of patterns and generalizations relating to sound inheritance reasoning; this reasoning was formalized in *TMOIS*, and metatheoretical results were established about the formalization. In this paper, we have described the discovery of several dimensions along which alternative approaches to inheritance reasoning can be generated, all of which appear to be equally sound.

The existence of these approaches has many theoretical consequences, which we are exploring in our current research. The inheritance theory of *TMOIS* needs to be generalized; and other treatments of nonmonotonic reasoning may need to be liberalized. For instance, the modal

nonmonotonic logic of [McDermott and Doyle 1980] renders the inference from $[A \wedge MB] \rightarrow B$ and A to $LB \vee L \sim B$ valid. On the natural interpretation of this inference in nets, it does hold good for credulous reasoners. But it fails for skeptical reasoners. This makes it desirable to search for a more general logic, that is neutral with respect to the design space of sound inheritance reasoners.

However, this theoretical work needs to be combined with more "empirical" investigations of the alternatives. Sandewall has suggested that a useful way to approach the question of sound multiple inheritance reasoning with exceptions proceed might be to create a catalog of problematic networks [Sandewall 1986]. Inheritance formalisms can then be compared by contrasting their handling of the cases in the catalog. As new definitions of inheritance are proposed, new examples may be added to the catalog to further refine the distinctions between the various schemes.

We have in effect been pursuing this approach. Although the various examples appearing throughout *TMOIS* were not collected in a formal catalog, the two primary examples, known as the Clyde level skip (figure 1) and the Nixon diamond (figure 2), were used in [Touretzky 1984] to contrast *TMOIS* with other formulations such as shortest path reasoning and Etherington's formalization of inheritance in Reiter's default logic [Etherington 1987]. Of course, the project of collecting examples has to be combined with extracting generalizations from the data and relating these generalizations to systematic theories. Thus, as well as presenting interesting examples where intuitions clash, we have presented in this paper a taxonomy of inheritance systems, and have presented four areas where definitions of multiple inheritance can disagree. A more systematic catalog is presented in [Touretzky *et al.* 1987].

Even after the careful analysis conducted in *TMOIS*, the problem of nonmonotonic multiple inheritance is far from conclusively resolved; its subtleties continue to surprise us.

Acknowledgements

We thank David Etherington and Erik Sandewall for illuminating correspondence, and Jon Doyle for helpful comments on an earlier draft of this paper.

References

- [Belnap 1977a] Belnap, N. D. How a computer should think. In G. Ryle (ed.), *Contemporary aspects of philosophy*. Oriol Press, 1977, pp. 30-56.
- [Belnap 1977b] Belnap, N. D. A useful four-valued logic. In J. M. Dunn and G. Epstein (eds.), *Modern uses of multiple-valued logic*, D. Reidel Publishing Company, 1977, pp. 8-37.
- [Borgida 1985] Borgida, A. Exceptions and taxonomies: knowledge engineering aspects of representation languages. Technical memo, Rutgers University Department of Computer Science, February, 1985.
- [Etherington 1987] Etherington, D. W. Formalizing non-monotonic reasoning systems. *Artificial Intelligence*, vol. 31, 1987, pp. 41-85.
- [Fahlman 1979] Fahlman, S. E. *NETL: A System for Representing and Using Real-World Knowledge*, The MIT Press, Cambridge, MA, 1979.
- [Fahlman *et al.* 1981] Fahlman, S. E., Touretzky, D. S., and van Roggen, W. Cancellation in a parallel semantic network. *Proceedings of IJCAI-81*, 1981, pp. 257-263.
- [Horty *et al.* 1987] Horty, J. F., Thomason, R. H., and Touretzky, D. S. A skeptical theory of inheritance in non-monotonic semantic nets. *Proceedings of AAAI-87*, 1987. A more complete version is available as a Carnegie Mellon University Computer Science Department technical report.
- [McDermott and Doyle 1980] McDermott, D. and Doyle, J. Non-monotonic logic I. *Artificial Intelligence*, vol. 13, 1980, pp. 41-72.
- [Reiter 1980] Reiter, R. A logic for default reasoning. *Artificial Intelligence*, vol. 13, 1980, pp. 81-132.
- [Roberts and Goldstein 1977] Roberts, R. B. and Goldstein, I. P. The FRL manual. AI Memo No. 409, MIT Artificial Intelligence Laboratory, Cambridge, MA, 1977.
- [Sandewall 1986] Sandewall, E. Non-monotonic inference rules for multiple inheritance with exceptions. *Proceedings of the IEEE*, vol. 74, 1986, pp. 1345-1353.
- [Thomason *et al.* 1986] Thomason, R. H., Horty, J. F., and Touretzky, D. S. A calculus for inheritance in monotonic semantic nets. Technical Report CMU-CS-138, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, 1986.
- [Touretzky 1984] Touretzky, D. S. Implicit ordering of defaults in inheritance systems. *Proceedings of AAAI-84*, 1984, pp. 322-325.
- [Touretzky 1986] Touretzky, D. S. *The Mathematics of Inheritance Systems*. Morgan Kaufmann Publishers, Los Altos, CA, 1986.
- [Touretzky *et al.* 1987] Touretzky, D. S., Horty, J. F., and Thomason, R. H. Issues in the design of nonmonotonic inheritance systems. Forthcoming technical report, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA, 1987.