

Mixing Strict and Defeasible Inheritance

John F. Horty^{1,2}
Richmond H. Thomason³

¹Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

²Philosophy Department
University of Maryland
College Park, MD 20742

³Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260

Abstract: Commonsense or expert knowledge of any rich domain involves an intricate mixture of strict and defeasible information. The importance of representing defeasible information in an inheritance system has been widely recognized, but it is not enough for a system to represent *only* defeasible information: without the ability to represent strict information as well, the system cannot represent definitional relations among concepts. As a response to this difficulty, we present a well-defined and intuitively attractive theory of inheritance for IS-A hierarchies containing strict and defeasible links types mixed together.

1 Introduction

It is apparent that commonsense or expert knowledge of any reasonably rich domain has to involve an intricate mixture of strict and defeasible information. The importance of representing defeasible information in a knowledge base—particularly, in a frame- or network-based inheritance reasoner—has been widely recognized. Although several early systems designed to allow defeasible inheritance reasoning were subject to semantic difficulties in their treatment of cancellation, these problems by now are essentially solved; there exist today a number of well-defined and intuitively attractive theories of defeasible inheritance. It has been argued, however, by Brachman [2] and Israel [6], that this intense concern with defeasible inheritance has obscured some of the more general problems exhibited by network formalisms capable of representing only defeasible information: without the ability to represent strict information as well, an inheritance system cannot express the important analytic or definitional relations among concepts.

One way of responding to these difficulties, exemplified by work in the KL-ONE tradition [3], is to abandon the project of encoding defeasible information in an inheritance reasoner, concentrating instead on definitional relations among richly structured concepts. As an alternative strategy, we are exploring ways in which strict and defeasible taxonomic information can be combined in a single inheritance network. This paper presents a first step: a well-defined and intuitively attractive theory of inheritance for IS-A hierarchies containing strict and defeasible link types mixed together. The analysis of mixed inheritance described here is itself a mixture, combining the

theory of strict inheritance from Thomason et al. [7] with the theory of defeasible inheritance provided by Horty et al. [5].

2 Basic concepts

Letters from the beginning of the alphabet (*a* through *d*) refer only to objects or individuals; letters from the middle of the alphabet (*m* through *t*) refer only to properties or kinds. Letters from the end of the alphabet (*u* through *z*) range over both objects and properties.

Where *y* is a property, the link types $x \Rightarrow y$ and $x \nRightarrow y$ represent positive and negative *strict* relations. If *x* is itself a property, these positive and negative strict links are equivalent to certain quantified conditionals: the link $p \Rightarrow q$ represents a statement of the form ‘Every *P* is a *Q*’; the link $p \nRightarrow q$ represents a statement of the form ‘No *P* is a *Q*’. If *x* is an object, these positive and negative strict links are equivalent to atomic and neg-atomic statements from ordinary logic: $a \Rightarrow p$ and $a \nRightarrow p$ represent the statements *Pa* and $\neg Pa$.

Where both *x* and *y* are properties, the link types $x \rightarrow y$ and $x \nrightarrow y$ represent *defeasible* relations. These positive and negative defeasible links are equivalent to ordinary generic statements: $p \rightarrow q$ and $r \nrightarrow q$, for example, might represent the statements ‘Birds fly’ and ‘Mammals don’t fly’. There is nothing in classical logic very close in meaning to generic statements like these. In particular, ‘Birds fly’ doesn’t mean that all birds fly, since it is true even in the presence of exceptions. Instead, it seems to mean that “typical birds” fly. Likewise, ‘Mammals don’t fly’ does not mean that no mammals fly, but only that “typical mammals” don’t.

Capital Greek letters represent *networks*—finite graphs, with nodes and link types as described. Networks are themselves classified as *strict* if they contain only strict links, or *defeasible* if they contain no strict links emanating from property nodes. *Mixed* networks can contain both strict and defeasible links emanating from property nodes.

Lower case Greek letters refer to *paths*—special sequences of links. Often, it is convenient to refer to an arbitrary path in a way that displays some of the nodes it passes through without displaying the particular link types connecting those nodes. For this purpose, we adopt a notation according to which ‘ $\pi(x, \sigma, y)$ ’ refers to an arbitrary positive path, and ‘ $\bar{\pi}(x, \sigma, y)$ ’ likewise to an arbitrary

negative path, from x through σ to y . As a convention governing this π -notation, we assume that adjacency of node symbols entails adjacency of nodes on the paths symbolized. Thus, for example, ' $\pi(x, u, \sigma, y)$ ' refers to a negative path beginning with a direct link of any type from x to u , and then moving through σ to y .

Paths are classified as simple or compound, strict or defeasible, positive or negative. The simple paths are just the direct links—classified as strict or defeasible, positive or negative, along with the links themselves. Compound paths are defined inductively, as follows. (1) If $\pi(x, \sigma, p)$ is a strict positive path, then: $\pi(x, \sigma, p) \Rightarrow q$ is a strict positive path; $\pi(x, \sigma, p) \nRightarrow q$ is a strict negative path; $\pi(x, \sigma, p) \rightarrow q$ is a defeasible positive path; and $\pi(x, \sigma, p) \nrightarrow q$ is a defeasible negative path. (2) If $\bar{\pi}(x, \sigma, p)$ is a strict negative path, then: $\pi(x, \sigma, p) \Leftarrow q$ is a strict negative path. (3) If $\pi(x, \sigma, p)$ is a defeasible positive path, then: $\pi(x, \sigma, p) \Rightarrow q$ is a defeasible positive path; $\pi(x, \sigma, p) \nRightarrow q$ is a defeasible negative path; $\pi(x, \sigma, p) \rightarrow q$ is a defeasible positive path; and $\pi(x, \sigma, p) \nrightarrow q$ is a defeasible negative path. (4) If $\bar{\pi}(x, \sigma, p)$ is a defeasible negative path, then: $\pi(x, \sigma, p) \Leftarrow q$ is a defeasible negative path.

It follows from this definition that an individual can occur in a path only as its initial node. Let us define a *negative segment* as a strict or defeasible negative link, possibly followed by a reverse positive strict path—that is, as a link sequence either of the form $x_1 \nrightarrow x_2 \Leftarrow \dots \Leftarrow x_n$ or of the form $x_1 \nRightarrow x_2 \Leftarrow \dots \Leftarrow x_n$. Then it follows from this definition also that if a negative segment occurs in a path, it can occur only at the very end.

Intuitively, paths represent arguments, which *enable* certain statements as their conclusions. A positive path of the form $\pi(x, \sigma, y)$ enables the statement $x \Rightarrow y$ if it is strict or x is an individual, and the statement $x \rightarrow y$ if it is defeasible and x is a kind; likewise, a negative path of the form $\bar{\pi}(x, \sigma, y)$ enables $x \nRightarrow y$ if it is strict or x is an individual, and $x \nrightarrow y$ if it is defeasible and x is a kind. Given a network Γ , the purpose of an inheritance theory is to specify the set of statements *supported* by Γ —that is, the set of statements we can reasonably conclude from the statements contained in Γ . We arrive at this specification in a roundabout way, defining a statement as supported by Γ just in case it is enabled by a path that Γ permits. It remains only to define the paths *permitted* by Γ —intuitively, the arguments sanctioned in the context of Γ .

3 Motivation

Since our approach to mixed inheritance combines the theory of strict inheritance from [7] with the theory of defeasible inheritance from [5], we first summarize these two theories, and then explain how they fit together.

For strict networks, our definition of permitted paths is very simple. According to the theory of [7], a strict network Γ permits exactly the paths it contains—that is, Γ permits σ iff σ is a path entirely composed of links con-

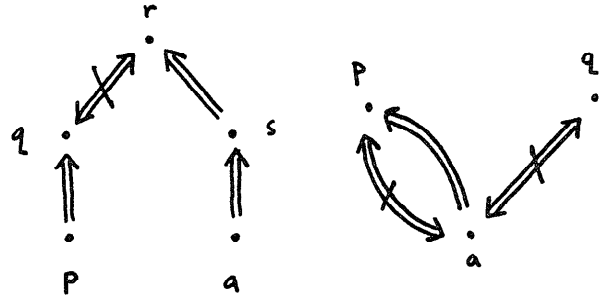


Figure 1: Γ_1

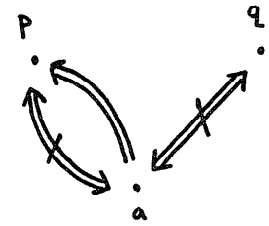


Figure 2: Γ_2

tained in Γ . In the case of Γ_1 (Figure 1), for example, the permitted paths include $a \Rightarrow s \Rightarrow r$ and $p \Rightarrow q \nRightarrow r \Leftarrow s$. Suppose we interpret the nodes in this net so that p = starlings, q = birds, r = mammals, s = dogs, and a = Rover. Then the first of these paths shows us how Γ_1 supports the conclusion that Rover is a mammal ($a \Rightarrow r$); the second shows how it supports the conclusion that no starlings are dogs ($p \nRightarrow s$). The net does not permit, for example, the path $p \Rightarrow s \Rightarrow r$, since the link $p \Rightarrow s$ is not contained in Γ_1 .

It is important to note that this analysis of strict inheritance is not the standard view. Strict networks contain only strict links, each of which is equivalent to a formula of classical logic. It may seem natural, then, to use classical logic itself to provide a semantics for such a network—by identifying the network with the set of formulas that translate its links, and then defining a statement as supported by the network just in case it belongs to the deductive closure of that set. This idea, which we take to be the standard view, is due originally to Hayes [4]. To see that it is different from the theory of [7], consider, for example, the net Γ_2 (Figure 2). This network would be translated into the set $\{Pa, \neg Pa, \neg Qa\}$. Since the set is inconsistent, any statement at all belongs to its classical deductive closure; so according to the standard view, Γ_2 should be taken to support every statement—including, say, Qa . According to the analysis of [7], however, Γ_2 does not support Qa , since it permits no positive path from a to q , and in fact provides uncontested evidence that $\neg Qa$.

It is, in some ways, a delicate matter to decide between the analysis of strict inheritance provided by [7] and the traditional analysis of [4]. One is always free to regard a strict network simply as a notational variant of some classical theory, so that the analysis of [4] would be appropriate. Still, there seems to be some value in taking seriously the graph-based nature of inheritance reasoners, which derive conclusions corresponding only to actual paths. The problem is then to see whether we can make logical sense of such a reasoner by designing an appropriate logic, rather than by forcing the reasoner to conform to the standards of an already-existing logic. This task is carried out for strict networks in [7], which provides both a Gentzen-style proof theory for path-based inheritance reasoning and an interpretation of the resulting logic in a four-valued model

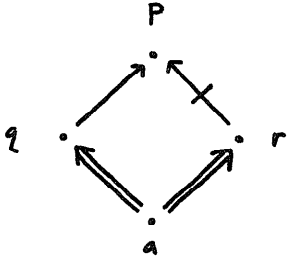


Figure 3: Γ_3

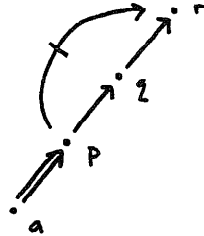


Figure 4: Γ_4

based on that of Belnap [1].

Defeasible inheritance is more complicated than strict inheritance, primarily because defeasible networks, unlike their strict counterparts, do not permit all the paths they contain. The distinction derives from the different role played in the two kinds of networks by conflicting paths. Any strict network containing conflicting paths is inconsistent, though—as the theory of [7] shows—the effects of the inconsistency can be localized. But defeasible networks can contain conflicting paths without even local inconsistency. Consider, for example, Γ_3 (Figure 3). Although this net contains conflicting paths, it is not inconsistent: it admits interpretations under which all of its links represent true statements—including the well-known interpretation with a = Nixon, q = Quakers, r = Republicans, and p = pacifists. Since the net is consistent, no reasonable theory of defeasible inheritance would say that it permits both of the conflicting paths $a \Rightarrow q \rightarrow p$ and $a \Rightarrow r \not\rightarrow p$ at once. Any such theory would allow us to draw inconsistent conclusions—for example, that Nixon both is a pacifist ($a \Rightarrow p$) and that he isn't ($a \not\Rightarrow p$)—from consistent information.

Theories of defeasible inheritance differ among themselves in their treatment of conflicting paths. One kind of theory associates with each network containing conflicting paths a number of different extensions, corresponding to different resolutions of the conflicts. Because each extension supports a maximal set of conclusions (subject to certain constraints) we describe these theories as *credulous*; an example is the theory of Touretzky [8]. The present paper is based on an alternative approach to defeasible inheritance, developed in [5], which has the advantage of associating with any given network only a single extension. We describe this alternative as a *skeptical* approach—since it embodies the broadly skeptical idea that conflicting arguments, represented in networks by conflicting paths, tend to neutralize each other. Applied to Γ_3 , for example, the skeptical approach tells us that neither of the conflicting paths should be permitted.

In fact, the theory of [5] is not thoroughly skeptical: its skepticism is restricted to compound paths, and even compound paths can be neutralized only by those conflicting paths that are not themselves, as we say, preempted. The first of these restrictions has the effect that, even in

the face of conflicts, any non-compound path contained in a network—that is, any direct link—will be permitted by that network. As explained in [5], this principle is well-motivated, particularly against the background of the four-valued logic; but it is not a crucial feature of the theory.

The second restriction, however, is crucial. Preemption is the mechanism by which, in case of conflicts, arguments based on more specific information are allowed to override arguments based on less specific information. For example, consider Γ_4 (Figure 4)—with, say, a = Tweety, p = penguins, q = birds, and r = flying things. Since this net contains the two conflicting paths $a \Rightarrow p \rightarrow q \rightarrow r$ and $a \Rightarrow p \not\rightarrow r$, an unrestricted skepticism would permit neither. However, it seems in this case that the latter of these paths *should* be permitted, because it represents an argument based on more specific information. The second restriction above reflects this intuition. We say that a path of the form $\pi(x, \tau, v) \rightarrow y$ is *preempted* in a net Γ just in case there is a node z such that (i) Γ permits a path $\pi(x, \tau_1, z, \tau_2, v)$, so that z provides “more specific” information than v about x , and (ii) $z \not\rightarrow y \in \Gamma$, so that z gives us “direct” information contrary to that provided by v . (The definition of preemption for negative paths is symmetrical.) According to the theory of [5], even a conflicted path will be permitted if the only paths with which it conflicts are themselves preempted; so, for example, $a \Rightarrow p \not\rightarrow r$ will be permitted by Γ_4 , since $a \Rightarrow p \rightarrow q \rightarrow r$ is preempted.

The theory of [7] tells us, then, that a strict network permits exactly the paths it contains; the theory of [5] tells us that a defeasible network permits a path it contains just in case that path is either a direct link, or any other path with which it conflicts is itself preempted. Now, to combine these two theories into an account of inheritance for mixed networks, we first carry over entirely the analysis of strict inheritance from [7], and then modify the analysis of defeasible inheritance from [5] to accommodate the presence of strict links. Since it incorporates the analysis of [7], the resulting theory tells us that a mixed network permits exactly the strict paths it contains. Likewise, since it is based on the analysis of [5], the resulting theory also embodies the skeptical idea that a compound defeasible path is neutralized by any conflicting path that is not itself preempted. However, in order to develop this idea in a mixed context, we need to modify slightly our conception of the kind of paths that represent conflicts, as well as our understanding of the preemption relation among conflicting paths.

In defeasible networks, all conflicts share a simple form: they involve paths with identical initial nodes, identical end nodes, and opposite polarity. But the presence of strict links introduces the possibility of less direct conflicts, even among defeasible paths. As an illustration, consider Γ_5 (Figure 5). Here it seems reasonable, in light of the strict segment $r \Rightarrow s \Rightarrow t$, to regard $p \rightarrow q \rightarrow r$ and $p \rightarrow u \rightarrow v \not\rightarrow t$ themselves as conflicting paths, even though they

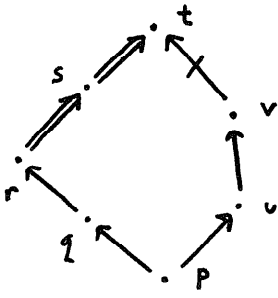


Figure 5: Γ_5

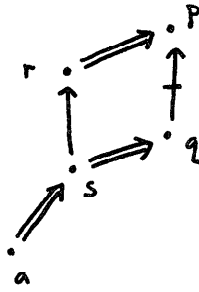


Figure 6: Γ_6

do not share an end node. Imagine, for example, that r = dogs, s = mammals, and t = animals, so that the strict segment tells us that *all* dogs are animals. In the context of Γ_5 , then, the path $p \rightarrow q \rightarrow r$, which represents an argument to the effect that p 's are dogs, carries with equal force the conclusion that p 's are animals; so it conflicts with $p \rightarrow u \rightarrow v \rightarrow t$, which represents an argument that p 's are not animals.

What this example shows is that two defeasible paths can represent conflicting arguments, even if they have different end nodes, when one of the paths clashes with a strict consequence of the other. Of course, such strict consequences can themselves be classified as positive or negative. Let us define $\kappa_\Gamma(x) = \{x\} \cup \{y : \Gamma \text{ contains a strict positive path from } x \text{ to } y\}$ and $\bar{\kappa}_\Gamma(x) = \{y : \Gamma \text{ contains a strict negative path from } x \text{ to } y\}$, so that $\kappa_\Gamma(x)$ and $\bar{\kappa}_\Gamma(x)$ represent the positive and negative strict consequences attributed to x by Γ —the set of properties that x must possess, according to Γ , and the set of properties that x cannot possess. It is then natural to extend our conception of conflicting defeasible paths so that, in addition to the ordinary kinds of clashes, a path of the form $\pi(x, \sigma, u) \rightarrow y$ will be said to conflict in a net Γ with any path of the form $\pi(x, \tau, v) \not\rightarrow m$ where $m \in \kappa_\Gamma(y)$, and also with any path of the form $\pi(x, \tau, v) \rightarrow m$ where $m \in \bar{\kappa}_\Gamma(y)$. Our general skeptical attitude regarding conflicting paths will then have to apply to these new kinds of conflicts as well. In Γ_5 , for instance, neither $p \rightarrow q \rightarrow r$ nor $p \rightarrow u \rightarrow v \rightarrow t$ will be permitted, since each is neutralized by its conflict with the other.

Just as the presence of strict links allows for the possibility of new kinds of conflicts, however, it provides also for the possibility of new relations of preemption. To see this, consider the network Γ_6 (Figure 6), supplied with an interpretation under which a = Hermann, p = persons born in America, q = native speakers of German, r = persons born in Pennsylvania, and s = native speakers of Pennsylvania Dutch. Under this interpretation, Γ_6 tells us is that Hermann is a particular speaker of Pennsylvania Dutch, that every speaker of Pennsylvania Dutch speaks German (since Pennsylvania Dutch is a dialect of German), that German speakers tend not to be born in America, that speakers of Pennsylvania Dutch tend to be born in Pennsylvania, and that everyone born in Pennsylvania is born

in America.

According to our new, extended conception, the paths $a \Rightarrow s \rightarrow r$ and $a \Rightarrow s \Rightarrow q \not\rightarrow p$ now represent conflicting arguments in the context of Γ_6 , since $p \in \kappa_{\Gamma_6}(r)$. Of course, we would not want to remain skeptical in this case. The path $a \Rightarrow s \Rightarrow q \not\rightarrow p$, representing the argument that Hermann was not born in America since he is a native speaker of German, should be *preempted* in Γ_6 : the fact that his dialect is Pennsylvania Dutch provides a more specific argument to the contrary. Without modification, however, the treatment of preemption from [5] does not give us this result. A path can be preempted only if there is more specific and direct information to the contrary; and, although s does provide “more specific” information than q , the path $s \rightarrow r \Rightarrow p$ does not represent “direct” information to the contrary—at least, not according to the standards of [5], which holds that direct information can be carried only by direct links.

Evidently, it is this last requirement concerning the nature of direct information that needs to be modified in the present context. In the context of defeasible networks, it makes good sense to say that direct information can be carried only by direct links: any compound path represents an argument that can itself be undermined. In the context of mixed nets, however, certain kinds of compound paths *can* legitimately be thought to carry direct information—namely, compound paths consisting of a single defeasible link followed by a strict end segment, of any length. In Γ_6 , for example, the path $s \rightarrow r \Rightarrow p$ should be thought of as telling us directly that speakers of Pennsylvania Dutch are born in America: for even by the standards of [5], $s \rightarrow r$ counts as a direct statement of the fact that speakers of Pennsylvania Dutch are born in Pennsylvania, and the strict extension $r \Rightarrow p$ simply tells us that *everyone* born in Pennsylvania is born in America.

Adjusting our definition of preemption to account for this new notion of direct information, we say now that a negative path $\pi(x, \tau, v) \not\rightarrow m$ is preempted in a mixed network Γ if there exist nodes z and n such that Γ permits a path $\pi(x, \tau_1, z, \tau_2, v)$ with $z \rightarrow n \in \Gamma$ and $m \in \kappa_\Gamma(n)$. This new definition allows us to conclude, as it should, that $a \Rightarrow s \Rightarrow q \not\rightarrow p$ is preempted in Γ_6 ; so the net does end up supporting the conclusion that Hermann was born in America. It is a bit more complicated to formulate mixed preemption for positive paths, although no new ideas are involved, simply because direct information to the contrary can now take the form either of a positive defeasible link followed by a negative strict extension, or of a negative defeasible link followed by a reverse positive strict extension. Formally, we say that a positive path $\pi(x, \tau, v) \rightarrow m$ is preempted in a mixed network Γ if there exist nodes z and n such that Γ permits a path $\pi(x, \tau_1, z, \tau_2, v)$ with either (i) $z \rightarrow n \in \Gamma$ and $m \in \bar{\kappa}_\Gamma(n)$ or (ii) $z \not\rightarrow n \in \Gamma$ and $n \in \kappa_\Gamma(m)$.

4 The definition

In this section, we assemble our motivational ideas into a definition of the permission relation for mixed networks; we use the symbol ' \vdash ' to stand for the permission relation, so that ' $\Gamma \vdash \sigma$ ' means that the net Γ permits the path σ . Like that of [5], the present definition is inductive. Our first step, then, is to assign a measure of "complexity" to each path σ in a net Γ in such a way that it can be decided whether $\Gamma \vdash \sigma$ once it is known whether $\Gamma \vdash \sigma'$ for each path σ' less complex in Γ than σ itself.

In order to arrive at the appropriate notion of path complexity, we proceed through a number of auxiliary ideas. As we recall, a path is a joined sequence of links containing a negative segment, if at all, only at the very end. Let us say, then, that a *generalized path* is a sequence of links joined like an ordinary path, except that it can contain negative segments anywhere, and perhaps more than one. (Example: $p \nrightarrow q \leftarrow r \nrightarrow s \leftarrow t$ is a generalized path, but it is not a path, since its negative segment $p \nrightarrow q \leftarrow r$ is not an end segment.) Next, let us define the *defeasible length* of a generalized path as follows: if the generalized path does not contain a strict initial segment, then its defeasible length is simply the number of defeasible links in the path; if the generalized path does contain a strict initial segment, then its defeasible length is the number of defeasible links in the path augmented by one. (Example: the generalized path $r \rightarrow s \Rightarrow t \rightarrow u$ has a defeasible length of two, since it contains two defeasible links and no strict initial segment; the generalized path $p \Rightarrow q \Rightarrow r \rightarrow s \Rightarrow t \rightarrow u$ is three, since it contains a strict initial segment along with two defeasible links.)

Using these ideas, we can now define the *degree* of a path σ in a net Γ —written, $\deg_{\Gamma}(\sigma)$ —as the greatest defeasible length of any acyclic generalized path in Γ from the initial node of σ to its end node. (Example: $\deg_{\Gamma_5}(p \rightarrow q \rightarrow r) = 3$, since the acyclic generalized path from p to r in Γ_5 whose defeasible length is greatest is $p \rightarrow u \rightarrow v \nrightarrow t \leftarrow s \leftarrow r$, with a defeasible length of 3.) In order to insure that the assignment of degree to the paths in a network has the appropriate properties, we need to restrict the application of our theory, as in [5], to paths free from certain kinds of defeasible cycles (a defeasible cycle is a cyclic generalized path containing at least one defeasible link); for the present, we limit will our attention, even more severely than necessary, to networks which are either entirely acyclic, or which contain only strict cycles.

The notion of degree defined here is a straightforward generalization of the notion defined in [5]. However, it is not quite appropriate as a measure of path complexity for an inductive definition of the permission relation; in the present context, the measure of complexity needs to carry just a bit more information. Basically, we want our measure of a path's complexity to tell us, in addition its degree, whether or not the path possesses a strict end segment. Therefore, we define the *mixed degree* of a path σ in a net Γ as a pair $\langle n, v \rangle$. The first component of the pair tells us the degree of σ in Γ : $n = \deg_{\Gamma}(\sigma)$. The second component

tells us, simply, whether or not σ possesses a strict end segment: by convention, we let $v = 0$ if σ does not possess a strict end segment, and $v = 1$ if it does. We define a lexical ordering on the mixed degrees by giving priority to the first component: $\langle n, v \rangle < \langle n', v' \rangle$ iff either $n < n'$ or $n = n'$ and $v < v'$. The idea behind this ordering is that degree is the primary measure of path complexity—but of two paths identical in degree, one with and one without a strict end segment, the path lacking the strict end segment is classified as less complex.

Our definition of the permission relation has the overall structure of a definition by cases. Any path σ from a mixed network can be divided into the subpaths $\mu(\sigma)$ and $\delta(\sigma)$, where $\mu(\sigma)$ is the maximal strict end segment of σ , and $\delta(\sigma)$ is the result of truncating $\mu(\sigma)$ from σ . (Example: if σ is $x \Rightarrow y \rightarrow p \nrightarrow r \leftarrow s$, then $\mu(\sigma)$ is $p \nrightarrow r \leftarrow s$ and $\delta(\sigma)$ is $x \Rightarrow y \rightarrow p$.) Using this notation, then, we specify the conditions under which $\Gamma \vdash \sigma$ in three separate cases, depending on the form of σ . Our first case deals with defeasible paths possessing strict end segments.

Case A: $\sigma \neq \delta(\sigma)$ and $\sigma \neq \mu(\sigma)$. Then $\Gamma \vdash \sigma$ iff $\Gamma \vdash \delta(\sigma)$ and $\Gamma \vdash \mu(\sigma)$.

The next case deals with strict paths.

Case B: $\sigma = \mu(\sigma)$. Then $\Gamma \vdash \sigma$ iff each link in σ is contained in Γ .

Finally, we deal with the case of paths ending in defeasible links—which itself divides into subcases, as such paths may be simple or compound.

Case C-I: $\sigma = \delta(\sigma)$ and σ is a direct link. Then $\Gamma \vdash \sigma$ iff $\sigma \in \Gamma$.

Case C-II: $\sigma = \delta(\sigma)$ and σ is a compound path. Two subcases to consider.

1. σ is a positive path, of the form $\pi(x, \sigma_1, u) \rightarrow y$. Then $\Gamma \vdash \sigma$ iff
 - (a) $\Gamma \vdash \pi(x, \sigma_1, u)$;
 - (b) $u \rightarrow y \in \Gamma$;
 - (c) For $m \in \kappa_{\Gamma}(y)$, $x \nrightarrow m \notin \Gamma$ and $m \notin \bar{\kappa}_{\Gamma}(x)$;
 - (d) For $m \in \bar{\kappa}_{\Gamma}(y)$, $x \rightarrow m \notin \Gamma$ and $m \notin \kappa_{\Gamma}(x)$;
 - (e) For all v, m, τ such that $\Gamma \vdash \pi(x, \tau, v)$ with $v \nrightarrow m \in \Gamma$ and $m \in \kappa_{\Gamma}(y)$, there exist z, n, τ_1, τ_2 such that $\Gamma \vdash \pi(x, \tau_1, z, \tau_2, v)$ with $z \rightarrow n \in \Gamma$ and $m \in \kappa_{\Gamma}(n)$;
 - (f) For all v, m, τ such that $\Gamma \vdash \pi(x, \tau, v)$ with $v \rightarrow m \in \Gamma$ and $m \in \bar{\kappa}_{\Gamma}(y)$, there exist z, n, τ_1, τ_2 such that $\Gamma \vdash \pi(x, \tau_1, z, \tau_2, v)$ with either (i) $z \rightarrow n \in \Gamma$ and $m \in \bar{\kappa}_{\Gamma}(n)$ or (ii) $z \nrightarrow n \in \Gamma$ and $n \in \kappa_{\Gamma}(m)$.
2. σ is a negative path, of the form $\pi(x, \sigma_1, u) \nrightarrow y$. Then $\Gamma \vdash \sigma$ iff

- (a) $\Gamma \vdash \pi(x, \sigma_1, u)$;
- (b) $u \not\vdash y \in \Gamma$;
- (c) For m such that $y \in \kappa_\Gamma(m)$, $x \rightarrow m \notin \Gamma$ and $m \notin \kappa_\Gamma(x)$;
- (d) For all v, m, τ , such that $\Gamma \vdash \pi(x, \tau, v)$ with $v \rightarrow m \in \Gamma$ and $y \in \kappa_\Gamma(m)$, there exist z, n, τ_1, τ_2 such that $\Gamma \vdash \pi(x, \tau_1, z, \tau_2, v)$ with either (i) $z \rightarrow n \in \Gamma$ and $m \in \kappa_\Gamma(n)$ or (ii) $z \not\vdash n \in \Gamma$ and $n \in \kappa_\Gamma(m)$.

It should be clear that this definition, although structured as a definition by cases, is properly an induction on mixed degree. Case A defines permission for a path σ of mixed degree $\langle n, 1 \rangle$ in terms of the path $\delta(\sigma)$ of mixed degree $\langle n, 0 \rangle$ and the path $\mu(\sigma)$ of mixed degree $\langle 1, 1 \rangle$ —both inductively simpler. Cases B and C-I are basis cases, defining permission respectively for paths of mixed degree $\langle 1, 1 \rangle$ and $\langle 1, 0 \rangle$. Finally, Case C-II defines permission for paths of mixed degree $\langle n, 0 \rangle$ with $n > 1$ in terms of paths of mixed degree $\langle n', v' \rangle$ —where v' may be either 0 or 1, but $n' < n$ so that the overall measure of mixed degree is simpler.

5 Conclusion

By combining the analysis of strict inheritance from [7] with the skeptical analysis of defeasible inheritance from [5], we have developed a well-defined and intuitively attractive theory of inheritance for semantic networks containing both strict and defeasible links. At this point, two topics stand out as the most important areas for further research.

The first concerns the treatment of cyclic networks. A central advantage of the theories of defeasible inheritance presented in both [5] and [8] is the ease with which they handle relations of preemption among conflicting arguments, naturally preferring those arguments based on more specific information. Part of what makes this possible is the restriction of these theories to acyclic networks, which allows us to define a partial ordering of “specificity” among the various argument paths. In the purely defeasible case, it is not terribly unnatural to restrict ourselves to acyclic networks; however, it is almost impossible to introduce strict links into a network without also introducing cyclic generalized paths. This paper limits itself to networks whose only cycles are entirely strict, but that limitation seems excessive. We need to discover the extent to which cyclic paths can be admitted into inheritance networks without destroying the partial ordering of specificity among arguments that makes a natural treatment of preemption possible.

The second research topic concerns the application of this work to the representation of complex concepts, such as Brachman’s “yellow elephant” or the traditional “unmarried man.” In order to represent such concepts along with defeasible information in a taxonomic reasoner, it is necessary, first, to develop a theory of inheritance allowing

for the expression of both strict and defeasible relations. The present paper presents such a theory—but it does not address the problem of handling complex concepts within the framework it sets out. A central accomplishment of the KL-ONE tradition has been the design and analysis of algorithms for handling complex defined concepts in an strict inheritance network. It is important, now, to begin exploring the degree to which this accomplishment can be duplicated in the context of a mixed inheritance reasoner.

Acknowledgments

This material is based on work supported by the National Science Foundation under Grant No. IRI-8700705. We thank Dave Touretzky for helpful discussions.

References

- [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] R. Brachman. “I lied about the trees” or, defaults and definitions in knowledge representation. *The AI Magazine*, vol. 6 (Fall, 1985), pp. 80–93.
- [3] R. Brachman and J. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, vol. 9 (1985), pp. 171–216.
- [4] P. Hayes. The logic of frames. In *Frame Conceptions and Text Understanding*, D. Metzger (ed.), Walter de Gruyter and Co. (1979), pp. 46–61. Reprinted in *Readings in Knowledge Representation*, R. Brachman and H. Levesque (eds.), Morgan Kaufmann (1985), pp. 287–297.
- [5] J. Horty, R. Thomason, and D. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. In *Proceedings of AAAI-87*, Morgan Kaufmann (1987), pp. 358–363.
- [6] D. Israel and R. Brachman. Some remarks on the semantics of representation languages. In *On Conceptual Modeling: Perspectives from Artificial Intelligence, Databases, and Programming Languages*, M. Brodie, J. Mylopoulos, and J. Schmidt (eds.), Springer-Verlag (1984).
- [7] R. Thomason, J. Horty, and D. Touretzky. A calculus for inheritance in monotonic semantic nets. In *Proceedings of the Second International Symposium on Methodologies for Intelligent Systems*, Z. Ras and M. Zemankova (eds.), North-Holland (1987), pp. 280–287.
- [8] D. Touretzky. *The Mathematics of Inheritance Systems*. Morgan Kaufmann (1986).