

# Skeptical Inheritance: Computing the Intersection of Credulous Extensions

Lynn Andrea Stein\*

Department of Computer Science

Brown University

Providence, RI 02912

## Abstract

*Ideally skeptical inheritance* supports exactly those inferences true in every credulous extension of an inheritance hierarchy. We provide a formal definition of ideally skeptical inheritance. We show that two path-based approaches fail to capture ideally skeptical inheritance, and that there are inheritance hierarchies for which there are more always-true inferences than always-supported paths. We describe an ATMS-like scheme that computes ideally skeptical inheritance and represents hierarchical dependencies using a limited form of Boolean satisfiability. Finally, we demonstrate a preemption (specificity) strategy for which ideally skeptical inheritance is polynomial-time computable.

## 1 Introduction

Inheritance, like other forms of defeasible reasoning, sanctions uncertain conclusions. In unambiguous contexts, these conclusions are only as reliable as the general rules from which they are derived: concluding that a particular bird can fly is reasonable only if birds generally fly. Ambiguity introduces another kind of uncertainty: if we have evidence that Charlie is a bird, and other evidence that he is not, our reasoning is still more uncertain because our assumptions are questionable.




Figure 1: A modified "Nixon Diamond."

Credulous reasoning involves this second kind of uncertainty. A credulous reasoner sanctions any internally consistent state of the world. For example, in figure 1, credulous inheritance allows either the conclusion that Nixon is a pacifist, or that Nixon is not a pacifist. In either case, the choice is arbitrary. Once we accept one

of these conclusions, however, we can reason further: if we assume that Nixon is a pacifist, we can conclude that, like a typical pacifist, he is anti-military. Credulous inheritance permits us to believe that Nixon *is not* a pacifist, or *is* a pacifist, and therefore *is* anti-military, but *either way our beliefs must be internally consistent*.

While credulous inheritance offers many alternatives, skeptical inheritance yields a single, unambiguous set of conclusions for any inheritance hierarchy. *Ideally skeptical inheritance* supports exactly those conclusions true in every credulous extension. In the next section, we provide formal definitions of ambiguity, credulous extensions, and ideally skeptical inheritance. We then present two approximate approaches to skeptical inheritance, ambiguity blocking and ambiguity propagating inheritance. The failure of these path-based approaches to be both sound and complete for ideally skeptical inheritance indicates the importance of reasoning about conclusions, or inferences, rather than about their supporting paths, or arguments. In section 5, we demonstrate an ATMS-like labeling scheme for inheritance hierarchies, based on a limited form of Boolean satisfiability. This labeling gives a precise description of dependencies in the hierarchy, including an exact representation of ideally skeptical inheritance. Finally, in section 6, we present a definition of preemption, or specificity, and provide a polynomial-time algorithm for computing the results of specificity. Although the definition of specificity itself is neither skeptical nor credulous, the output of the specificity algorithm can be used as input to the ideally skeptical inheritance algorithm to yield a polynomial-time algorithm for computing *ideally skeptical inheritance with specificity*.

## 2 Ambiguity and credulous extensions

An *inheritance hierarchy*  $\Gamma = \langle V_\Gamma, E_\Gamma \rangle$  is a directed acyclic graph with positive and negative edges, intended to denote "is-a" and "is-not-a" respectively. We write a positive edge from  $a$  to  $x$  as  $a \cdot x$ , and a negative edge  $a \cdot \neg x$ . We call a sequence of positive edges  $a \cdot s_1 \cdots s_n \cdot x$  ( $n \geq 0$ ) a *positive path*, and a sequence of positive edges followed by a single negative edge  $a \cdot s_1 \cdots s_n \cdot \neg x$  ( $n \geq 0$ ) a *negative path*.<sup>1</sup> We use lower case Greek letters to stand

<sup>1</sup>The approach described in this paper is upwards reasoning. That is, inheritance works from the focus node up-

\*The author is supported by an IBM Graduate Fellowship.

for sequences of positive edges, so these positive and negative paths might be abbreviated  $a \cdot \sigma \cdot x$  and  $a \cdot \sigma \cdot \neg x$ . We takes  $*$  to be a variable ranging over the set  $\{., \neg\}$ :  $s * x$  stands for any edge from  $s$  to  $x$  ( $s \cdot x$  or  $s \cdot \neg x$ ) and  $s \nrightarrow x$  to stand for the edge of opposite sign ( $s \cdot \neg x$  or  $s \cdot x$ , respectively).

A path, or *argument*,  $a \cdot \sigma * x$  supports the inference “ $a$  is (not) an  $x$ .” We use the notation  $a \rightarrow x$  ( $a \nrightarrow x$ ) to stand for this inference, or conclusion, independently of the path through which it is derived. One inference—e.g.,  $a \rightarrow x$ —may have many supporting arguments— $a \cdot \sigma \cdot x$ ,  $a \cdot \tau \cdot x$ , etc. We use  $\sim$  in much the same way as  $*$ , as a variable ranging over the set  $\{\rightarrow, \nrightarrow\}$ , and say that  $a \sim x$  if either  $a \rightarrow x$  or  $a \nrightarrow x$ . Throughout much of this paper, paths and inferences will be used interchangeably; in section 4, however, the distinction will become important.

An inheritance hierarchy  $\Gamma$  supports a path  $a \cdot \sigma * x$ , written  $\Gamma \triangleright a \cdot \sigma * x$ , if the path  $a \cdot s_1 \cdots s_n * x$  is in  $E_\Gamma$  and it is admissible;  $\Gamma$  supports an inference  $a \sim x$  if it supports some corresponding path. Initially, we will take “admissible” to be vacuous: any path actually in  $\Gamma$  is supported. This is credulous inheritance at its most general, so that any conflicting paths are potentially ambiguous and result in corresponding credulous extensions. In section 6, we extend this definition to include a specificity criterion analogous to Touretzky’s [1986] inferential distance.

A brief aside on semantics is relevant here. If every conflict is taken to be ambiguous, preemption strategies may be viewed as ambiguity-resolving heuristics—preferences as to how these ambiguities should be resolved. This results in a preference-based semantics for inheritance, where credulous extensions play a role analogous to models in Shoham’s [1988] model-preference semantics for nonmonotonic logics. This is precisely the view we are espousing here; a more complete exploration of the approach may be found in [Stein, 1990].




Figure 2: Unambiguous w.r.t.  $a$ , ambiguous w.r.t.  $b$ .

Ambiguity arises when two supported paths conflict. Formally, an inheritance hierarchy  $\Gamma$  is *ambiguous* w.r.t. a node  $a$  if there is some node  $x \in V_\Gamma$  such that  $\Gamma \triangleright a \rightarrow x$  and  $\Gamma \triangleright a \nrightarrow x$ . In this case, we say that the ambiguity is *at*  $x$ . Ambiguity is always relative to a node: for example, the hierarchy in figure 2 is unambiguous w.r.t.  $a$ , but ambiguous w.r.t.  $b$  (at  $e$ ). Credulous inheritance is a means of resolving these ambiguities.

A *credulous extension* of an inheritance hierarchy  $\Gamma$  with respect to a node  $a$  is a maximal unambiguous subhierarchy of  $\Gamma$  with respect to  $a$ : if  $X^{\Gamma,a}$  is a credulous extension of  $\Gamma$  w.r.t.  $a$ , then there is no unambiguous subgraph of  $\Gamma$  supporting a proper superset of the paths

wards, rather than from the root downwards. The ramifications of upwards versus downwards reasoning are discussed in [Touretzky et al., 1987]; some complexity concerns are described in [Selman and Levesque, 1990].




Figure 3: A network and some extensions.

that  $X^{\Gamma,a}$  supports. An example of a network with several credulous extensions—and some non-extensions—is given in figure 3. If  $X^{\Gamma,a}$  is a credulous extension of  $\Gamma$  w.r.t.  $a$ ,  $a$  is called the *focus node* of  $X^{\Gamma,a}$ . If an extension  $X^{\Gamma,a} \triangleright a \rightarrow x$ ,  $x$  is *true* in  $X^{\Gamma,a}$ ; similarly  $a \nrightarrow x$  and *false* in  $X^{\Gamma,a}$ .

In general, there may be many credulous extensions of  $\Gamma$ . If  $\Gamma$  is ambiguous w.r.t.  $a$ , there will be several extensions of  $\Gamma$  w.r.t.  $a$ .  $\{X_i^{\Gamma,a}\}_i$  denotes the set of credulous extensions of  $\Gamma$  w.r.t.  $a$ . Further,  $\{X_i^{\Gamma,a}\}_i$ —the set of credulous extensions w.r.t.  $a$ —may differ from  $\{X_i^{\Gamma,b}\}_i$ : ambiguities which arise in determining what  $a$  is may never arise in determining what  $b$  is. However, computing the credulous extensions of a hierarchy w.r.t. multiple foci is independent of the order in which they are considered.<sup>2</sup>

Choosing a credulous extension involves making an arbitrary choice. *Skeptical* inheritance is intended to select only those inferences that are not arbitrary. These are precisely the inferences which hold in any credulous extension, or possible state of the world. Formally, an inheritance hierarchy  $\Gamma$  skeptically permits an inference  $a \sim x$ , written  $\Gamma \vdash a \sim x$ , if for every credulous extension  $X^{\Gamma,a} \in \{X_i^{\Gamma,a}\}_i$ ,  $X^{\Gamma,a} \triangleright a \sim x$ .  $\vdash$  defines *ideally skeptical inheritance*.

Unfortunately, there is no “skeptical extension” corresponding to the definition of ideally skeptical inheritance. The closest we can come is  $\Omega(\Gamma, a)$ , the subgraph of  $\Gamma$  containing those edges that are in every credulous extension of  $\Gamma$  w.r.t.  $a$ :<sup>3</sup>  $\Omega(\Gamma, a) = \langle V_\Gamma, \bigcap_{\{X_i^{\Gamma,a}\}_i} E_{X_i^{\Gamma,a}} \rangle$ .  $\Omega(\Gamma, a)$  supports exactly those arguments, or sequences of edges, that are in every credulous extension. However, there may be inferences  $a \sim x$  that are supported in every credulous

<sup>2</sup>It might therefore seem logical to define a credulous extension independently of any particular node; this would simply be a maximal subgraph with no ambiguity w.r.t. any node. Unfortunately, this is not terribly useful, as disambiguating one node (say,  $b$  in figure 2) can require eliminating perfectly good paths for another node (e.g. the path  $a \rightarrow e$ ).

<sup>3</sup>Throughout this paper, we use the notation  $\aleph(\Gamma, a)$ , where  $\aleph$  stands for upper case Greek letter, to refer to a uniquely determined subhierarchy of  $\Gamma$  w.r.t.  $a$ : for example, the skeptical extension  $\Omega(\Gamma, a)$ ; the ambiguity blocking extension  $B(\Gamma, a)$ ; the ambiguity propagating extension  $\Pi(\Gamma, a)$ ; the specificity extension  $\Sigma(\Gamma, a)$ .

extension of  $\Gamma$  w.r.t.  $a$ , but have different supporting arguments in different extensions: some extensions may contain one sequence of edges— $a \cdot \sigma \cdot x$ —while other extensions contain a different sequence of edges— $a \cdot \tau \cdot x$ . Hierarchies with this property are discussed in section 4. This means that,  $\Omega(\Gamma, a)$ —and any other path-based “skeptical extension”—can support at most a subset of the inferences ideally skeptical inheritance admits; capturing the complete set of inferences requires reasoning about cases.

### 3 Ambiguity blocking inheritance

The first attempt at skeptical inheritance was taken by Horty *et al.* [1987].<sup>4</sup> They argue that an ambiguous line of reasoning should not be allowed to interfere with other potential conclusions. Because this approach discontinues a line of reasoning as soon as an ambiguity has been reached, we refer to it as *ambiguity blocking inheritance*. Although Horty *et al.* describe one specific theory—including, e.g., a particular preemption strategy—ambiguity blocking inheritance defines a general approach:

Starting from the focus node  $a$ , if a node  $x$  is ambiguous w.r.t.  $a$  in  $\Gamma$ , eliminate all edges into and out of  $x$ . When the entire hierarchy has been scanned, the remaining edges form a new network,  $B(\Gamma, a)$ , which is unambiguous w.r.t.  $a$ . This is the *ambiguity blocking skeptical extension* of  $\Gamma$ ; ambiguity blocking inheritance concludes that a network  $\Gamma$  admits  $a \rightsquigarrow x$  exactly when  $B(\Gamma, a) \triangleright a \cdot \sigma \cdot x$ .




Figure 4: Applying ambiguity blocking inheritance to  $\Gamma$  w.r.t.  $a$  yields  $B(\Gamma, a)$ .

While ambiguity blocking inheritance seems reasonable, it results in some anomalous conclusions. Consider, for example, figure 4. Ambiguity blocking inheritance on  $\Gamma$  with focus node  $a$  determines that  $e$  is ambiguous w.r.t.  $a$ , so it eliminates all edges to and from  $e$ . In particular, it eliminates the edge  $e \cdot f$ , making  $f$  unambiguous w.r.t.  $a$ :  $B(\Gamma, a) \triangleright a \rightarrow f$ . This is certainly one possibility. But it is also possible that  $a \rightarrow e$ ; and if  $a \rightarrow e$ , it is ambiguous whether  $a \rightarrow f$ —that is,  $a$  might not be an  $f$ . It is certainly not safe to assume from the ambiguity at  $e$  that

<sup>4</sup>According to Horty (personal communication), a “skeptical” approach to inheritance is one which offers a unique, unambiguous set of conclusions for any inheritance hierarchy. This differs with our intuition that “skeptical” means “unwilling to believe uncertain conclusions.” In Horty’s view, computing the intersection of the credulous extensions is only one way to reason “skeptically.”




Figure 5: Ambiguity blocking inheritance computes “parity”:  $a$  is not an  $e$ , is a  $g$ , is not an  $h$ , is a  $j$ ,.... Also,  $a$  is a  $j$ ,  $b$  is not a  $j$ ,  $d$  is a  $j$ ,  $f$  is not a  $j$ ,....

the path  $a \cdot b \cdot d \cdot f$  is always true. But this is precisely what ambiguity blocking inheritance does.<sup>5</sup>

A more severe anomaly follows from this first. Ambiguity blocking inheritance computes a kind of “parity” on the number of ambiguities in a path. According to Horty *et al.*, the network in figure 5 is skeptical as to whether  $a$  is-a  $e$  or an  $i$ , but supports the conclusions that  $a$  is-a  $g$  and a  $j$ . Similarly, this net is skeptical about whether  $b$  or  $f$  is-a  $j$ , but allows the paths from  $a$  and  $d$  to  $j$ . More than the first anomaly, this result calls into question the intuitiveness of ambiguity blocking inheritance. In any case, ambiguity blocking inheritance is *promiscuous*: there are inferences  $a \rightsquigarrow x$  such that  $B(\Gamma, a) \triangleright a \rightsquigarrow x$  but  $\Gamma \not\models a \rightsquigarrow x$ .

### 4 Ambiguity propagating inheritance

Unlike ambiguity blocking inheritance, *ambiguity propagating inheritance* allows ambiguous lines of reasoning to proceed. An argument thus cannot be certain unless there are *no* counterarguments; in contrast, ambiguity blocking inheritance considers only *unambiguous* counterarguments. Like ambiguity blocking inheritance, ambiguity propagation defines a family of algorithms. One such algorithm is outlined here; a complete definition and discussion may be found in [Stein, 1989]. This algorithm computes  $\Pi(\Gamma, a)$ , the ambiguity propagating skeptical extension of  $\Gamma$  w.r.t.  $a$ , in time  $O(E_\Gamma)$ :

Starting from the focus node  $a$ , if a node  $x$  is ambiguous w.r.t.  $a$  in  $\Gamma$ , rather than eliminating all edges to and from  $x$ , retain  $x$  but mark it as ambiguous, and continue inheriting. Although paths to and from  $x$  will not be in  $\Pi(\Gamma, a)$ , they can still act as counterarguments and prevent other nodes from being unambiguous.  $\Pi(\Gamma, a)$  is the subgraph of  $\Gamma$  with those edges  $x \cdot y \in E_\Gamma$  such that  $x$  and  $y$  are both unambiguous w.r.t.  $a$ .

For example, the cascading ambiguities of figures 4 and 5, which gave ambiguity blocking inheritance dif-

<sup>5</sup>Horty *et al.* originally pointed out this difference between ambiguity blocking inheritance and ideally skeptical inheritance in [Horty *et al.*, 1987] and [Touretzky *et al.*, 1987].




Figure 6: Ambiguity propagating inheritance draws no conclusions about whether  $a$  is-a  $e$  or  $g$  or  $i$  or  $j$ .

ficulty, present no problem for ambiguity propagating inheritance. Figure 6 gives the ambiguity propagating extension for the network in figure 5.

In fact,  $\Pi(\Gamma, a)$  never supports extra paths.  $\Pi(\Gamma, a)$  is a subset of  $\Omega(\Gamma, a)$ , the “best approximation” to a skeptical extension, with the further property that  $\Pi(\Gamma, a)$  and  $\Omega(\Gamma, a)$  support  $a \rightsquigarrow x$  for exactly the same nodes  $x$ .  $\Pi(\Gamma, a)$  and  $\Omega(\Gamma, a)$  both support those arguments, or sequences of edges, w.r.t.  $a$  that are true in all credulous extensions of  $\Gamma$  w.r.t.  $a$ . But there are some inferences  $a \rightsquigarrow x$  that are supported by every credulous extension of  $\Gamma$  w.r.t.  $a$ , but have different supporting arguments in different extensions. In these circumstances, we need to reason about inferences rather than paths.




Figure 7: Whether a seedless grape vine is a fruit plant, or an arbor plant, it is certainly a plant!

For example, consider the hierarchy in figure 7.<sup>6</sup> Every credulous extension w.r.t. *seedless grape vine* supports the inference *seedless grape vine*  $\rightarrow$  *plant*, so ideally skeptical inheritance concludes that *seedless grape vines* are *plants*. Suppose, for example, that a *seedless grape vine* is a *fruit plant*; then it is a *plant*. Suppose that it is *not* a *fruit plant*; then it is unambiguously an *arbor plant*, and therefore a *plant*. In any state of the world, no matter how we resolve the ambiguities of the taxonomy, a *seedless grape vine* is a *plant*.

If we wish to determine what is true in all possible worlds, we cannot avoid this kind of reasoning. There are facts which are true in all credulous extensions, but which have no justification in the intersection of those extensions. This is why we cannot generate a “skeptical extension”—no *particular* set of edges of  $\Gamma$  from *seedless grape vine* to *plant* is in every credulous extension, so no such path can be in the “skeptical extension.” Thus every path-based approach to skeptical inheritance will always be either unsound or incomplete with respect to

<sup>6</sup> Apparently, Matt Ginsberg independently proposed a hierarchy with similar properties, in which *Nixon* is always *politically motivated*.

ideally skeptical inheritance. We can only compute the always-true inferences by, in effect, reasoning about all of the credulous extensions. Fortunately, in acyclic hierarchies, such reasoning is tractable.

## 5 Ideally skeptical inheritance

Ideally skeptical inheritance, applied to an inheritance hierarchy  $\Gamma$ , computes those inferences  $a \rightsquigarrow x$  that have some supporting argument in every credulous extension of  $\Gamma$  w.r.t.  $a$ . We compute ideally skeptical inheritance by keeping track of the conditions under which an inference  $a \rightsquigarrow x$  has support. If  $a \rightsquigarrow x$  always has support, then it is true in every credulous extension of  $\Gamma$ .

Let  $C_{\text{pos}}$  be the “positive children” of  $x$ —the nodes  $p_i \in V_\Gamma$  with positive edges  $p_i \cdot x \in E_\Gamma$ —and let  $C_{\text{neg}}$  be the “negative children” of  $x$ . Then if none of  $C_{\text{pos}}$  is true (i.e., not  $a \rightarrow p_i$ ),  $x$  cannot be true: it has no support at all. If at least one of  $C_{\text{pos}}$  is true, but none of  $C_{\text{neg}}$  is true, then  $x$  is true: its true positive children provide it with a supporting argument. If at least one of  $C_{\text{pos}}$  is true, but at least one of  $C_{\text{neg}}$  is true, too, then  $x$  is ambiguous: it has some support, but there is also a counterargument.

We can phrase this as a constraint satisfaction problem:

$a$ : The focus node is *always* true:  $\Gamma \vdash a \rightarrow a$ .

$\bigvee_i C_{\text{pos}_i} \supset \bar{x}$ :  $x$  cannot be true if none of its positive children is.

$((\bigvee_i C_{\text{pos}_i}) \wedge \bigvee_i C_{\text{neg}_i}) \supset x$ : If  $x$  has support but no counterargument,  $x$  is true.

There is no constraint corresponding to the third condition; if  $x$  is ambiguous, it is unconstrained.

If we conjoin these constraints, we get the biconditional

$$x \equiv (((\bigvee_i C_{\text{pos}_i}) \wedge (\bigvee_i C_{\text{neg}_i})) \vee ((\bigvee_i C_{\text{pos}_i}) \wedge (\bigvee_i C_{\text{neg}_i}) \wedge x))$$

This is the condition under which  $\Gamma \vdash a \rightarrow x$ . In the following algorithm for computing ideally skeptical inheritance, we label  $x$  with a formula corresponding to this constraint. The notation  $\llbracket x \rrbracket = [y]$  means that the propositional variable  $y$  is the label of the node named  $x$ . The usual rules w.r.t. connectives,  $\top$ , and  $\perp$ , apply.

The label of  $x$  describes exactly those conditions under which the focus node  $a$  is-a  $x$ . If  $x$  is labeled  $\top$ , it is always true— $\Gamma \vdash a \rightarrow x$ . If  $x$  is labeled  $\perp$ , its label can never be satisfied, so there is no credulous extension of  $\Gamma$  w.r.t.  $a$  that supports  $a \rightarrow x$ . If a node  $x$  is labeled  $\varphi$ ,  $x$  is true ( $a \rightarrow x$ ) in exactly those credulous extensions where  $\varphi$  is true. For example, in the labeling of figure 1,  $\llbracket \text{anti-military} \rrbracket = [\text{pacifist}]$ ; that is, *Nixon* is anti-military in exactly those extensions in which he is also a pacifist.

To determine the properties of node  $a$ :<sup>7</sup>

1. Label  $\llbracket a \rrbracket = [\top]$ .
2. Label all leaves (other than  $a$ , if  $a$  is a leaf)  $[\perp]$ .

<sup>7</sup>This algorithm computes the conditions under which  $a \rightarrow x$ . The condition under which  $a \not\rightarrow x$  is given by  $(\bigvee_i \llbracket C_{\text{neg}_i} \rrbracket) \wedge (\bigvee_i \llbracket C_{\text{pos}_i} \rrbracket)$ .




Figure 8:  $a$  is an  $x$  whenever the label of  $x$  has truth value  $T$ .

- For a generic node,  $x$ , with positive children  $C_{pos_i}$  and negative children  $C_{neg_j}$ , label

$$[x] = [((\bigvee_i [C_{pos_i}]) \wedge (\bigvee_j [C_{neg_j}]) \wedge T) \vee ((\bigvee_i [C_{pos_i}] \vee \perp) \wedge (\bigvee_j [C_{neg_j}] \vee \perp) \wedge x)]$$

Figure 8 gives the skeptical labeling of the hierarchy of figure 7.

This labeling corresponds roughly to an “ATMS labeling” of the hierarchy—the label of a node is its justification. We can use the labeling for incremental update of the hierarchy. Consider, for example, the hierarchy of figure 1. If we later discover that Nixon is a pacifist, the labeling automatically tells us that he is antimilitary as well (since  $[anti-military] = [pacifist]$ , and now  $[pacifist] = [T]$ ). In fact, we can incorporate all sorts of ambiguity resolving information—from specificity, to domain-specific knowledge, to updated beliefs—into this labeling simply by adding further constraints.

The complexity of this extended labeling algorithm, including further constraints, is unknown. Since it is a special case of boolean satisfiability, the problem may be  $\mathcal{NP}$ -hard. However, for the limited case of determining that a label is falsifiable—i.e. that there is some credulous extension that does not include the node—Kautz and Selman [1989] provide a polynomial algorithm. This means that we can compute the exact intersection of credulous extensions—ideally skeptical inheritance—in polynomial time.

## 6 Specificity

In the discussion of skeptical inheritance above, we assumed that all paths in the hierarchy were equally acceptable. In this section, we describe a *specificity criterion*, or preemption strategy, that makes choices among certain competing paths. The idea of a specificity criterion dates from [Etherington and Reiter, 1983] and [Touretzky, 1986]. Since then, many definitions of

specificity have appeared in the literature, but all operate on the same underlying principle: more specific information is likely to be more accurate. For example, in figure 9, information about *penguins* is more specific to *Tweety* than information about *birds*, so we can infer that *Tweety* does not fly.

We define specificity recursively. Certainly, if  $a * x \in E_\Gamma$ , then  $\Gamma$  supports  $a \sim x$ —that is, both the path, and the inference “ $a$  is (not) an  $x$ .” But what about compound paths?

Consider the path  $a \cdot \sigma * x$  ( $\sigma = s_1 \dots s_n, n \geq 0$ ).  $\Gamma$  supports  $a \cdot \sigma * x$  w.r.t.  $a$  if

- None of the edges of  $a \cdot \sigma$  ( $a \cdot s_1$  and  $s_i \cdot s_{i+1}$ , ( $1 \leq i < n$ )) is redundant in  $\Gamma$  w.r.t.  $a$ ,
- $\Gamma$  supports  $a \cdot \sigma$ , and
- $a \cdot \sigma * x$  is not preempted in  $\Gamma$  w.r.t.  $a$ .

More generally,  $\Gamma$  supports  $b \cdot \tau * w$  w.r.t.  $a$  if for every set of positive edges  $\sigma$ , whenever  $\Gamma$  supports  $a \cdot \sigma \cdot b$ ,  $\Gamma$  also supports  $a \cdot \sigma \cdot b \cdot \tau * w$ .

The path  $a \cdot \sigma * x$  is *preempted* in  $\Gamma$  w.r.t.  $a$  there is some direct intermediary  $s_i \in \sigma$ , with  $i$  strictly less than  $n$ , such that  $\Gamma$  supports the path  $a \cdot s_1 \dots s_i \neq x$ .

An edge  $b \cdot w$  is *redundant* in  $\Gamma$  w.r.t. focus node  $a$  if  $\Gamma$  supports some path  $b \cdot q_1 \dots q_n \cdot w$ ,  $n$  strictly greater than 0, w.r.t.  $a$ , and  $\Gamma$  does not support  $a \cdot \tau \cdot \neg q_i$ , for any  $0 \leq i \leq n$ , w.r.t.  $a$ .

Although the definitions of supported and redundant are mutually dependent, they are not circular. Because the hierarchy is acyclic, it can be ordered topologically, and the definition of support for a path from  $a$  to  $x$  depends only on the redundancy of nodes strictly topologically earlier than  $x$ .

The difficulties caused by redundant links were noted by Touretzky [1986]: in the second hierarchy ( $\Gamma'$ ) in figure 9, the edge from *king penguin* to *bird* is redundant—*king penguins* are typically *birds* even without that link. However, if that edge is not excluded, there will be a path *Tweety* · *king penguin* · *bird* · *flies* for which no intermediate node has an edge to  $\neg \text{flies}$ . Clearly, this is not the intended meaning here (or, indeed, in any network of this form, since the “penguin” node is always more specific than the “bird” node).

This definition of specificity is *on-path* and *upwards*. On-path means that a path is preempted only if one of its member nodes is involved in a counterargument. In contrast, some preemption strategies also allow a path to be preempted by a counterargument originating with a node *off* the path. Upwards inheritance reasons about the properties of a particular object, rather than the objects possessing a particular property. Our definition most closely resembles an upwards version of [Touretzky, 1986].

If we examine a hierarchy,  $\Gamma$ , from the perspective of a particular node,  $a$ , specificity provides a means of pruning the hierarchy—removing those edges that have been preempted. We call this subhierarchy  $\Sigma(\Gamma, a)$ , the *specificity extension* of  $\Gamma$  w.r.t.  $a$ . For example, the specificity extensions of the hierarchies in figure 9 w.r.t. *king penguin* are shown in figure 10. The definition of specificity above always yields a unique specificity extension




Figure 9: Information about penguins is more specific than information about birds, w.r.t. Tweety.




Figure 10: Specificity extensions for the hierarchies of figure 9 w.r.t. king penguin.

for a hierarchy w.r.t. a focus node. Figure 11 gives a polynomial-time algorithm for calculating the specificity extension of an inheritance hierarchy.

Not all of the ambiguities in an inheritance hierarchy are susceptible to specificity. For example, “diamond” ambiguities such as the Nixon diamond of figure 1 cannot be resolved using any preemption technique. For this reason,  $\Sigma(\Gamma, a)$  may still contain ambiguities, and may yield several credulous extensions w.r.t.  $a$ . These credulous extensions are the possible world-states consistent with the specificity criterion, and are a subset of the possible world-states delimited by  $\Gamma$ —the credulous extensions of  $\Gamma$  w.r.t.  $a$ . In particular, the credulous extensions of  $\Sigma(\Gamma, a)$  are the *preferred* extensions of  $\Gamma$  w.r.t.  $a$ , according to the preference induced by specificity:

Let  $X_1^{\Gamma, a}$  and  $X_2^{\Gamma, a}$  be two credulous extensions of an inheritance hierarchy  $\Gamma$  w.r.t. focus node  $a$ . Then specificity prefers  $X_1^{\Gamma, a}$  to  $X_2^{\Gamma, a}$  ( $X_1^{\Gamma, a} \preceq X_2^{\Gamma, a}$ ) if there are some node  $x$  and non-redundant, preempted path  $a \cdot \sigma \cdot x$  in  $\Gamma$  such that  $X_2^{\Gamma, a}$  supports  $a \cdot \sigma \cdot x$  and  $X_1^{\Gamma, a}$  does not, and  $X_1^{\Gamma, a}$  and  $X_2^{\Gamma, a}$  agree on all nodes that topologically precede  $x$  (according to *some* topological sort).<sup>8</sup>

## 7 Discussion

Hierarchies like figure 7 demonstrate that determining universal truths requires reasoning about inferences rather than paths. It follows that any purely path-based approach to inheritance must be either unsound or incomplete for ideally skeptical inheritance. Some additional bookkeeping mechanism, such as the “ATMS labeling” introduced here, must be added to inheritance systems that perform this type of reasoning.

The framework for skeptical inheritance, and the preference-based semantics of [Stein, 1990], are both independent of a particular definition of specificity. However, the tractability of the specificity criterion in section 6 makes it particularly useful here. Recent results [Kautz and Selman, 1989, Selman and Levesque, 1990] indicate the intractability of many basic inheritance problems. Our specificity criterion provides a tractable basis for both credulous and skeptical inheritance reasoning.

Proofs of several results described in this paper, and further exploration of these ideas, may be found in [Stein, 1989].

<sup>8</sup>We could insist that  $X_1^{\Gamma, a}$  and  $X_2^{\Gamma, a}$  agree on all nodes that precede  $x$  in *any* topological sort. It turns out that the minimal elements under this definition of preference are equivalent to the minimal elements under the definition given here.

## Compute-Specificity-Extension

Let  $\Gamma$  be an inheritance hierarchy.  
Sort the nodes of  $\Gamma$  topologically  
For each focus node  $a$   $\Sigma(\Gamma, a)$  will be the specificity  
 $\Sigma(\Gamma, a) := \Gamma$  ; extension of  $\Gamma$  w.r.t.  $a$ .  
For each node  $x$  reachable from  $a$ , in topological order ; from  $a$  to ...  
For each edge  $p \cdot x$ , in reverse topological order  
Let  $\Gamma^* = \Sigma(\Gamma, a) - \{q \mid q \not\prec x \in E_{\Sigma(\Gamma, a)}\}$   
If  $\Gamma^*$  no longer contains a positive path from  $a$  to  $p$ ,  
then remove the edge  $p \cdot x$  from  $\Sigma(\Gamma, a)$   
For each remaining positive edge  $p \cdot x$   
If  $\Sigma(\Gamma, a)$  contains a path  $p \cdot q_1 \cdots q_n \rightarrow x$ ,  $n \geq 1$   
such that there is no negative path from  $a$  to any of  $q_i$ ,  $x$ ,  
then remove the edge  $p \cdot x$  from  $\Sigma(\Gamma, a)$

Figure 11: A polynomial-time algorithm to compute  $\Sigma(\Gamma, a)$

## Acknowledgments

Discussions with Leora Morgenstern, David Touretzky, Jeff Horta, Henry Kautz, and the (individual and collective) members of Brown University AI group were invaluable in this work. The comments of the two *IJCAI* referees were extremely helpful as well.

## References

- [Etherington and Reiter, 1983] David Etherington and Raymond Reiter. On inheritance hierarchies with exceptions. In *AAAI*, pages 104–108, Washington, DC, 1983. Morgan Kaufmann Publishers, Inc.
- [Horta et al., 1987] John F. Horta, Richmond H. Thomason, and David S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. In *AAAI*, 1987.
- [Kautz and Selman, 1989] Henry Kautz and Bart Selman. Hard problems for simple default logics. In *First International Conference on Knowledge Representation and Reasoning*, Toronto, Ontario, 1989. Morgan Kaufmann Publishers, Inc.
- [Selman and Levesque, 1990] Bart Selman and Hector Levesque. The tractability of path-based inheritance. In *Formal Aspects of Semantic Networks*. Morgan Kaufmann Publishers, Inc., 1990.
- [Shoham, 1988] Yoav Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, 1988.
- [Stein, 1989] Lynn Andrea Stein. A preference-based approach to inheritance. Technical report, Brown University Department of Computer Science, Providence, RI, 1989.
- [Stein, 1990] Lynn Andrea Stein. Credulous extensions as possible worlds. In *Formal Aspects of Semantic Networks*. Morgan Kaufmann Publishers, Inc., 1990.
- [Touretzky et al., 1987] David S. Touretzky, John F. Horta, and Richmond H. Thomason. A clash of intuitions: The current state of nonmonotonic multiple inheritance systems. In *IJCAI*, 1987.
- [Touretzky, 1986] David Touretzky. *The Mathematical Theory of Inheritance*. Morgan Kaufmann, 1986.