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.

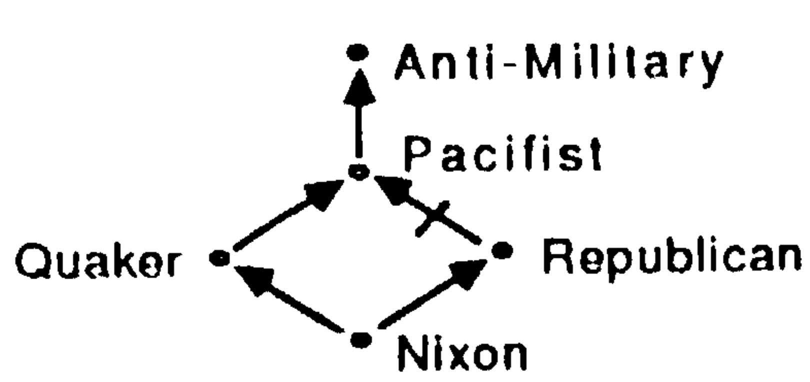

Figure 1: A modified "Nixon Diamond."

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* $T = (V_T, E_T)$ 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 $a.\neg x$. We call a sequence of positive edges $a.s_1 \ldots s_n.x$ ($n > 0$) a *positive path,* and a sequence of positive edges followed by a single negative edge $a.s_1 \ldots s_n \neg$ ($n > 0$) a *negative path}* We use lower case Greek letters to stand

*The author is supported by an IBM Graduate Fellowship.

*The approach described in this paper is upwards reasoning. That is, inheritance works from the focus node upfor sequences of positive edges, so these positive and negative paths might be abbreviated a α. x and a α . x. We takes * to be a variable ranging over the set {., .¬}: $s * x$ stands for any edge from s to x (s . x or s ¬x) and $s \not* x$ to stand for the edge of opposite sign (s .¬x or s. x, respectively).

A path, or *argument*, a.α*x supports the *inference* "a is (not) an x." We use the notation $a \rightarrow x$ $(a \not\rightarrow x)$ to stand for this inference, or conclusion, independently of the path through which it is derived. One inference—e.g., a—>x—may have many supporting arguments—a α.x, a . r . x, etc. We use ~> in much the same way as *, as a variable ranging over the set $\{\rightarrow, \not\rightarrow\}$, and say that a~>x if either a—>x or a-/+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 F *supports* a path a • a * x, written T \> a a * x, if the path as\ • • s_n*x is in Ep *and it is admissible:,* T supports an inference a~~>x if it supports some corresponding path. Initially, we will take "admissible" to be vacuous: any path actually in T 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^ [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 ShohanVs [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].

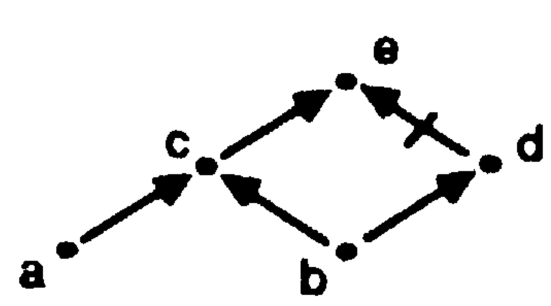

Figure 2: Unambiguous w.r.t. a, ambiguous w.r.t. b.

Ambiguity arises when two supported paths conflict. Formally, an inheritance hierarchy V is *ambiguous* w.r.t. a node a if there is some node $x \in V_\Gamma$ such that $\Gamma \rhd a \rightarrow x$ and $\Gamma \rhd a \not\rightarrow x$. 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 F with respect to a node a is a maximal unambiguous subhierarchy of F with respect to *a:* if $X^{\Gamma,a}$ is a credulous extension of F w.r.t. a, then there is no unambiguous subgraph of T supporting a proper superset of the paths 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 T w.r.t. a, a is called the *focus node* of $X^{\Gamma,a}$. If an extension $X^{\Gamma,a} \rhd a \rightarrow x$, x is *true in* $X^{\Gamma,a}$; similarly $a \not\rightarrow x$ and *false in* $X^{\Gamma,a}$.

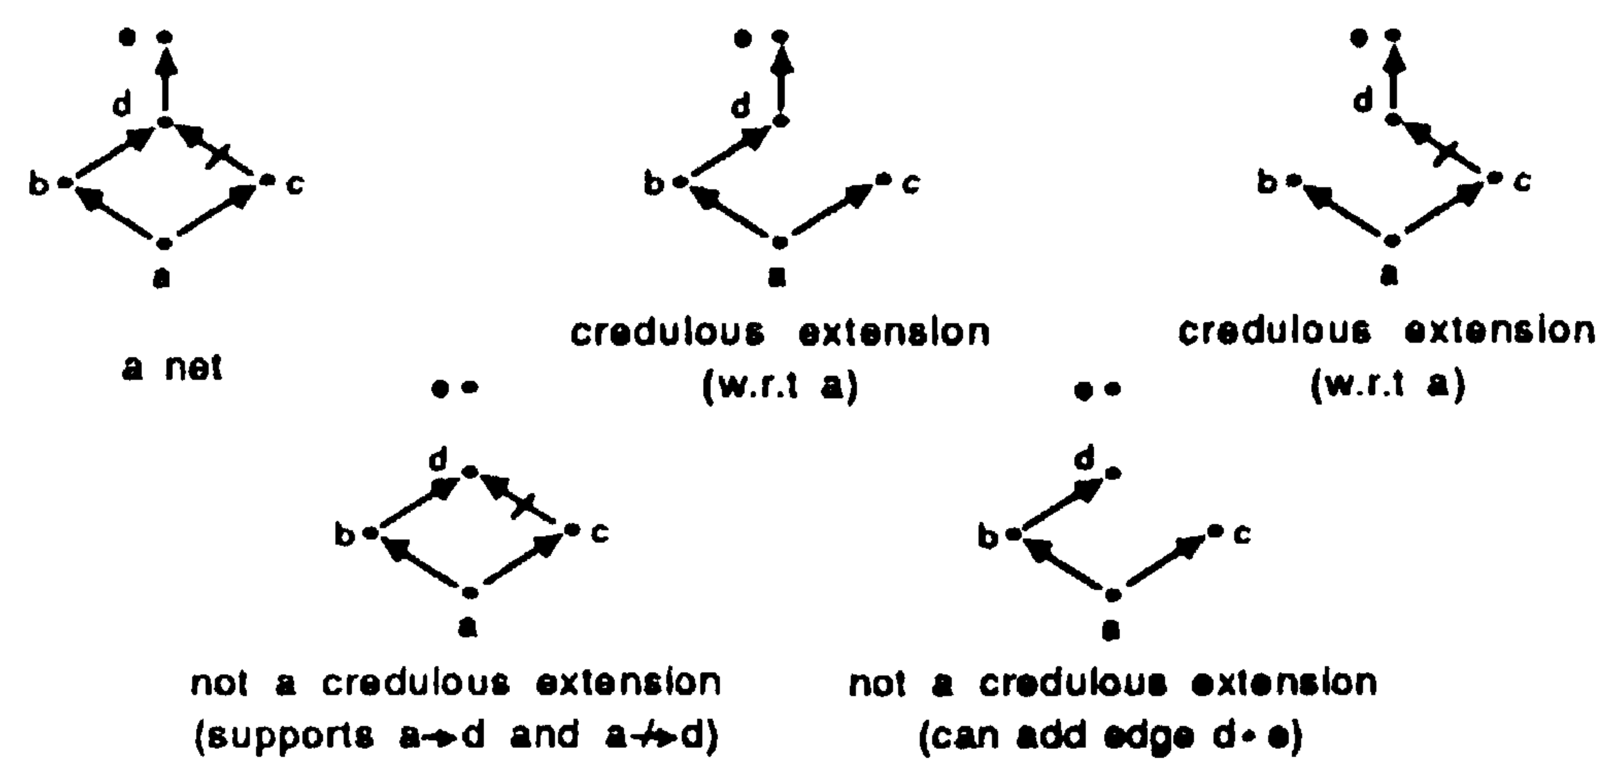

Figure 3: A network and some extensions.

In general, there may be many credulous extensions of T. If F is ambiguous w.r.t. a, there will be several extensions of F w.r.t. *a.* $\{X_i^{\Gamma,a}\}_i$ denotes the set of credulous extensions of F 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.[2]

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 F skeptically permits an inference $a \rightsquigarrow x$, written $\Gamma \mid\!\sim a \rightsquigarrow x$, if for every credulous extension $X^{\Gamma,a} \in \{X_i^{\Gamma,a}\}_i$, $X^{\Gamma,a} \rhd a \rightsquigarrow x$. $\mid\!\sim$ 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 T containing those edges that are in every credulous extension of F w.r.t. a:[3] $\Omega(\Gamma, a) = \langle V_\Gamma, \bigcap_{\{X^{\Gamma,a}\}_i} E_{X_i^{\Gamma,a}} \rangle$. $\Omega(\Gamma, a)$ supports exactly those arguments, or sequences of edges, that are in every credulous extension. However, there may be inferences $a \rightsquigarrow x$ that are supported in every credulous

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

[2]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*—>e).

[3]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 T w.r.t. o: for example, the skeptical extension $\Omega(\Gamma, a)$; the ambiguity blocking extension $\mathrm{B}(\Gamma, a)$; the ambiguity propagating extension n(P,a); the specificity extension $\Sigma(\Gamma, a)$.

extension of *T* 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].[4] 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 *ct. 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 I\ eliminate all edges into and out of *x.* When the entire hierarchy has been scanned, the remaining edges form a new network, B(r,a), which is unambiguous w.r.t. a. This is the *ambiguity blocking* skeptical extension of T; ambiguity blocking inheritance concludes that a network T admits a~>x exactly when B(T, a) |> a . *a* $*$ *x.*

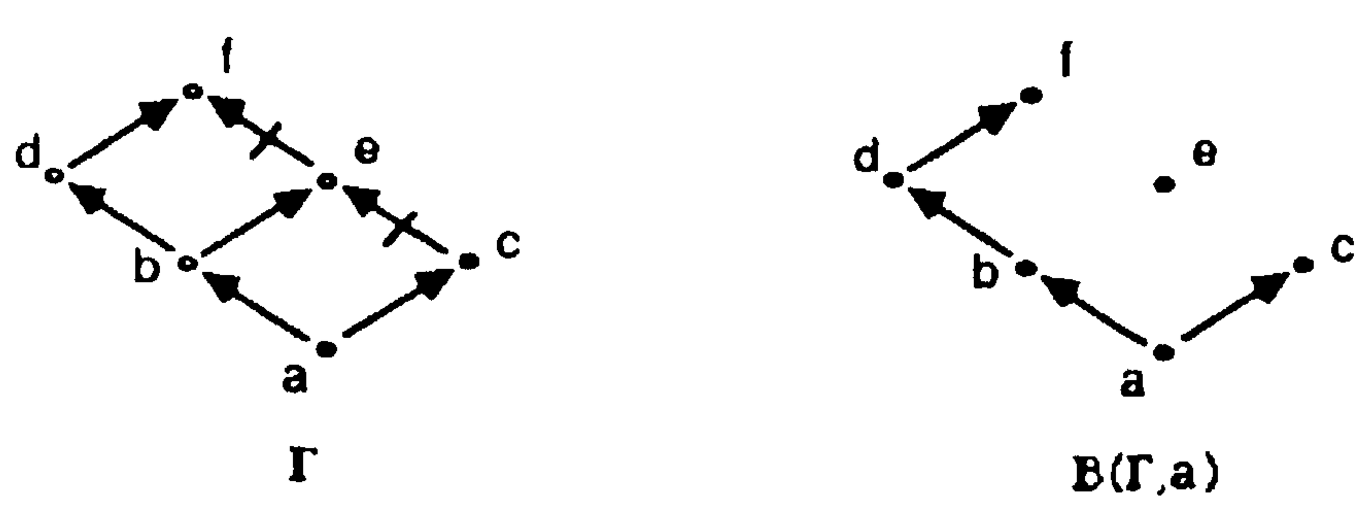

Figure 4: Applying ambiguity blocking inheritance to F w.r.t. *a* yields B(F,a).

While ambiguity blocking inheritance seems reasonable, it results in some anomalous conclusions. Consider, for example, figure 4. Ambiguity blocking inheritance on F 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* •/, making / unambiguous w.r.t. *a:* $B(\Gamma, a) \triangleright a \rightarrow f$. This is certainly one possibility. But it is also possible that a—>e; and if *a—>e,* it is ambiguous whether a—> f—that is, *a* might not be an /. It is certainly not safe to assume from the ambiguity at *e* that

[4]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."

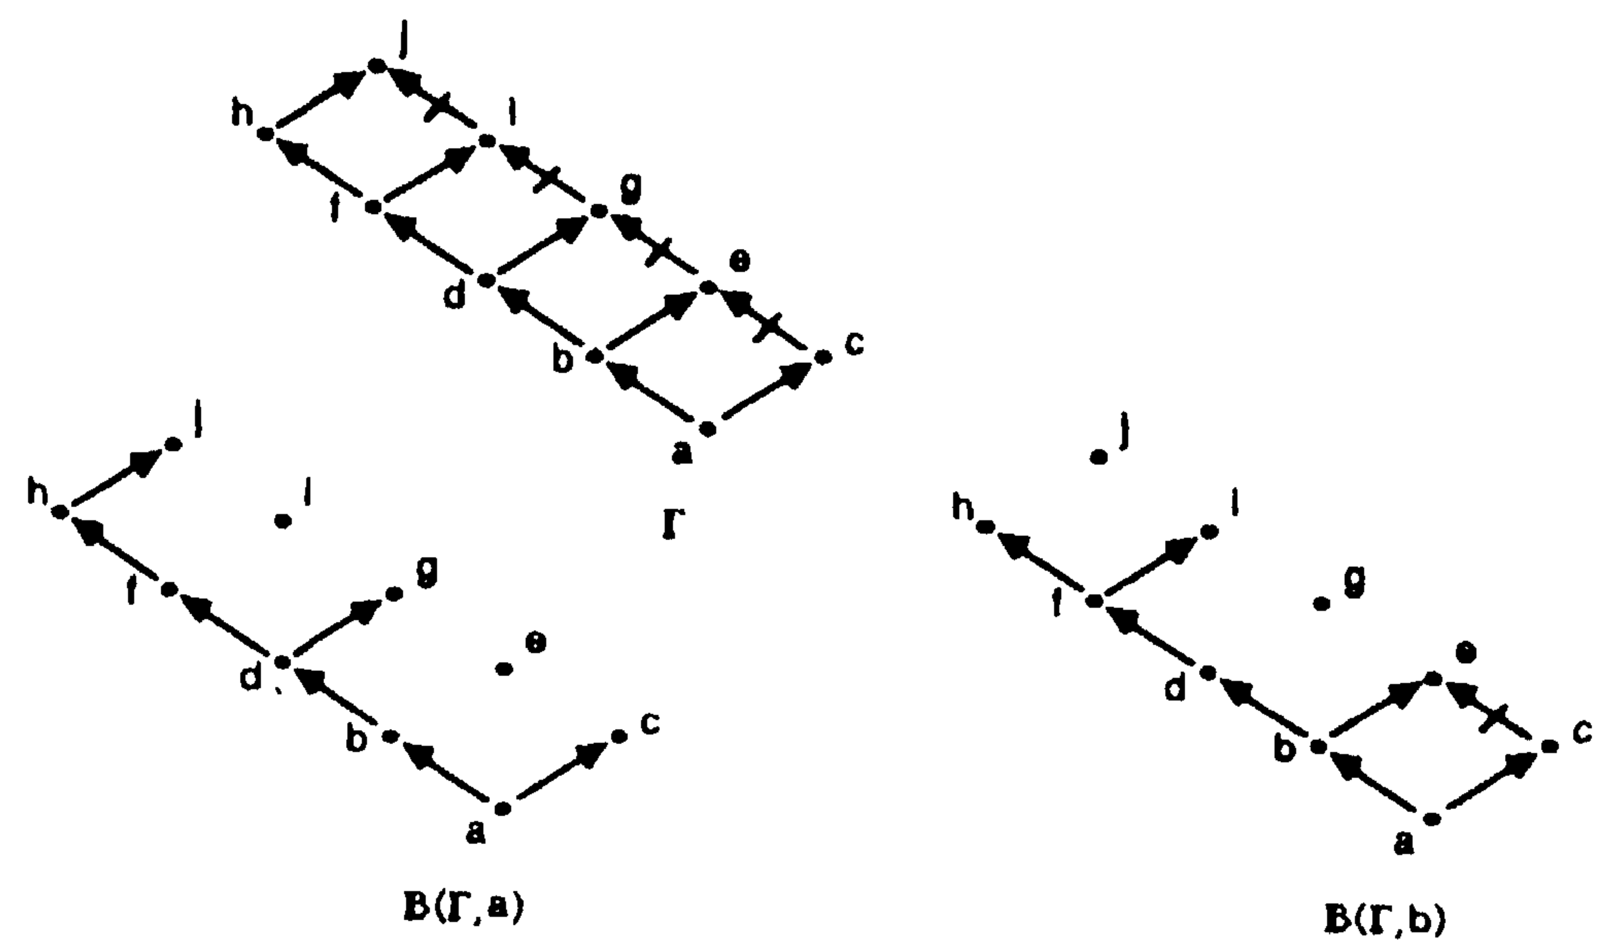

Figure 5: Ambiguity blocking inheritance computes "parity": a is not an e, is a g, is not an h, is a j,..... Also, a is a j, b is not a j, d is a j, f is not a j,.....

the path *a. b.d. f* is always true. But this is precisely what ambiguity blocking inheritance does.[5]

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 e*t. 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 / 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~>x* such that $B(\Gamma, a) \triangleright a \rightsquigarrow x$ but $\Gamma \not\triangleright a \rightsquigarrow x$.

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 F w.r.t. a, in time O(Ep):

> Starting from the focus node a, if a node *x* is ambiguous w.r.t. a in T, 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 u n a m b i g ι $\Pi(\Gamma, a)$ s the subgraph of T with those edges *x* * $y \in E_\Gamma$ 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-

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

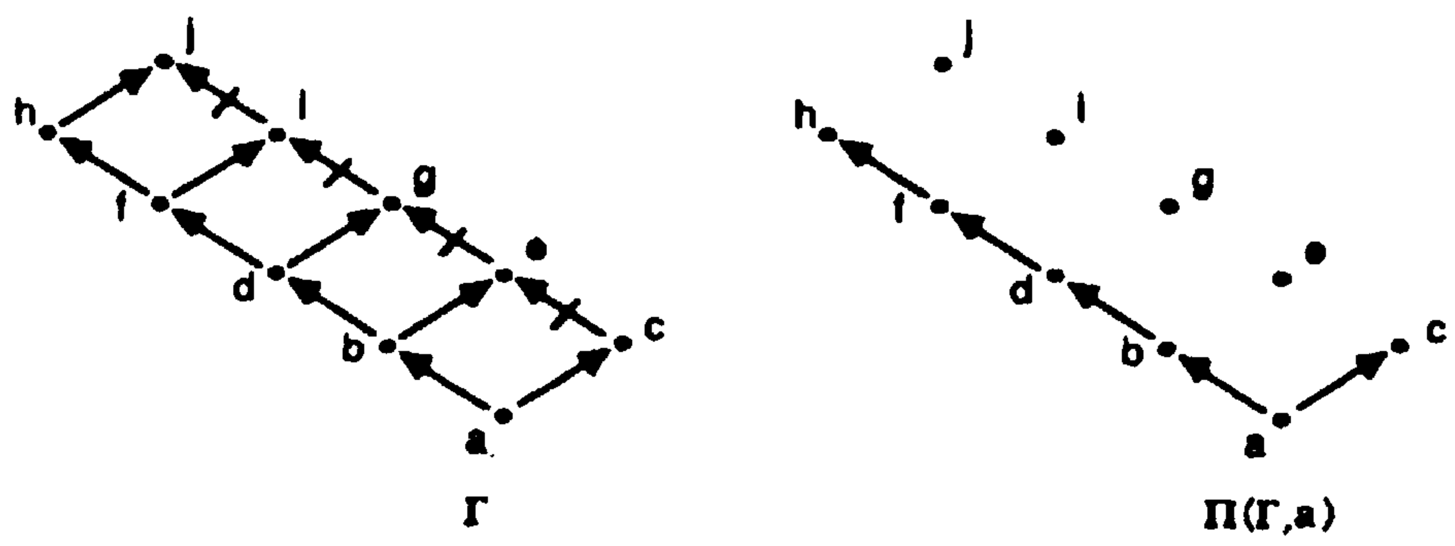

Figure 6: Ambiguity propagating inheritance draws no conclusions about whether *a* is-a *e* or *g* or *i* or *j*.

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 $a \leadsto x$ 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 Γ w.r.t. a. But there are some *inferences* $a \leadsto x$ that are supported by every credulous extension of Γ w.r.t. a, but have different supporting arguments in different extensions. In these circumstances, we need to reason about inferences rather than paths.

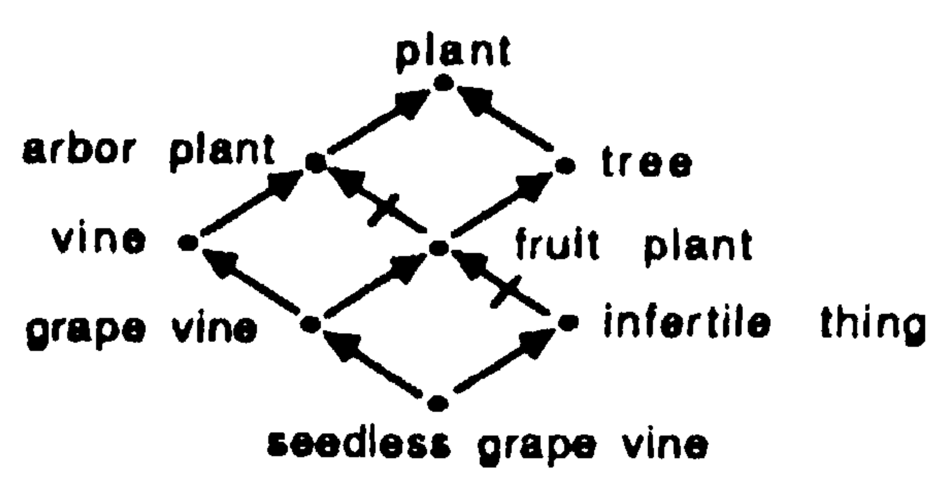

Figure 7: Whether a seedless grape vine is a fruit plant, or an arbor plant, it is certainly a plant!

For example, consider the hierarchy in figure 7.[6] Every credulous extension w.r.t. *seedless grape vine* supports the inference *seedless grape vine —> 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 *T* 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 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 Γ, computes those inferences $a \leadsto x$ that have some supporting argument in every credulous extension of Γ w.r.t. a. We compute ideally skeptical inheritance by keeping track of the conditions under which an inference $a \leadsto x$ has support. If $a \leadsto x$ always has support, then it is true in every credulous extension of Γ.

Let C_{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_{neg} be the "negative children" of x. Then if none of C_{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_{pos} is true, but none of C_{neg} is true, then x is true: its true positive children provide it with a supporting argument. If at least one of C_{pos} is true, but at least one of C_{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: $\Gamma \vdash a \rightarrow a$.

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

$((\vee_i C_{\text{pos}_i}) \wedge \overline{\vee_i C_{\text{neg}_i}}) \supset x$: 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

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

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 $[\![x]\!] = [y]$ means that the propositional variable y is the label of the node named x. The usual rules w.r.t. connectives, $\top$, and $\bot$, 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 $\bot$, its label can never be satisfied, so there is no credulous extension of Γ w.r.t. a that supports $a \rightarrow x$. If a node x is labeled φ, x is true ($a \rightarrow x$) in exactly those credulous extensions where φ is true. For example, in the labeling of figure 1, $[\![$*anti-military*$]\!]$ = [*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:[7]

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

[6] Apparently, Matt Ginsberg independently proposed a hierarchy with similar properties, in which *Nixon* is always *politically motivated.*

[7] This algorithm computes the conditions under which $a \rightarrow x$. The condition under which $a \not\rightarrow x$ is given by $(\vee_i [\![C_{\text{neg}_i}]\!]) \wedge (\overline{\vee_i [\![C_{\text{pos}_i}]\!]})$.

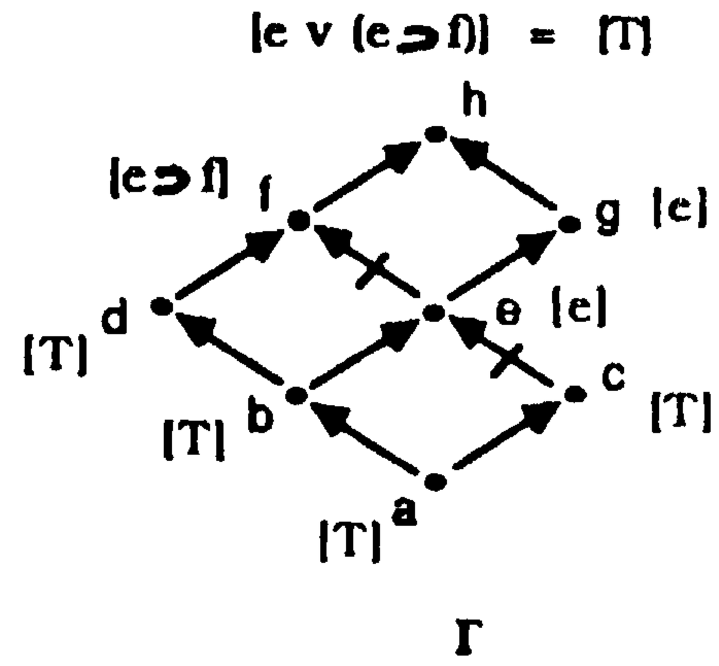

Figure 8: a is an x whenever the label of x has truth value ⊤.

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

$$[\![x]\!] = [((\bigvee_i [\![C_{pos_i}]\!]) \wedge (\overline{\bigvee_i [\![C_{neg_i}]\!]}) \wedge \top) \vee ((\bigvee_i [\![C_{pos_i}]\!] \vee \bot) \wedge (\bigvee_i [\![C_{neg_i}]\!] \vee \bot) \wedge x)]$$

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 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.

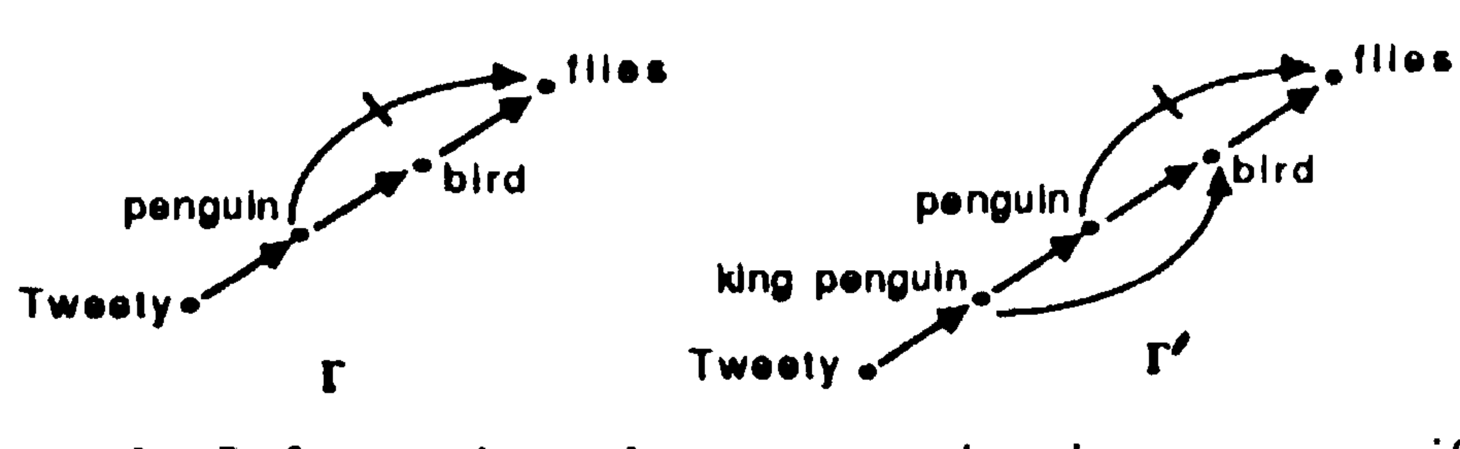

Figure 9: Information about penguins is more specific than information about birds, w.r.t. Tweety.

We define specificity recursively. Certainly, if ⊠T,* x ∈ then T supports $a \leadsto x$—that is, both the path, and the inference "a is (not) an *x*" But what about compound paths?

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

1. 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 Γ w.r.t. a,
2. Γ supports $a \cdot \sigma$, and
3. $a \cdot \sigma * x$ is not preempted in Γ w.r.t. a.

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

The path $a \cdot \sigma * x$ is *preempted* in Γ w.r.t. a there is some direct intermediary $s_i \in \sigma$, with i strictly less than n, such that Γ supports the path $a \cdot s_1 \cdots s_i \not* x$.

An edge $b \cdot w$ is *redundant* in Γ w.r.t. focus node a if Γ supports some path $b \cdot q_1 \cdots q_n \cdot w$, n strictly greater than 0, w.r.t. a, and Γ 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 (T') 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 *¬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 dafinition most closely resembles an upwards version of [Touretzky, 1986].

If we examine a hierarchy, T, 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 F 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

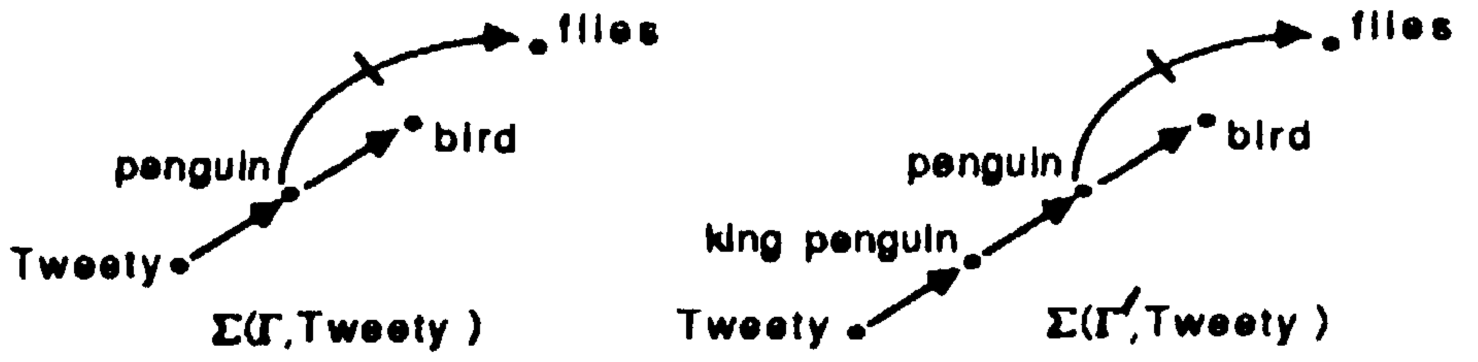

Figure 10: Specificity extensions for the hierarchies of figure 9 w.r.t. *king penguin.*

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, Σ(T,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 *T*—the credulous extensions of *T* w.r.t. *a.* In particular, the credulous extensions of E(T, a) are the *-preferred* extensions of T 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 Γ 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 $a \cdot \sigma * x$ in Γ such that $X_2^{\Gamma,a}$ supports $a \cdot \sigma * x$ 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).[8]**

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].

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 Γ be an inheritance hierarchy.
Sort the nodes of Γ topologically
For each focus node a                          ; Σ(Γ, a) will be the specificity
  Σ(Γ, a) := Γ                                 ; extension of Γ w.r.t. a.
  For each node x reachable from a, in topological order        ; from a to ...
    For each edge p • x, in reverse topological order
      Let Γ* = Σ(Γ, a) − {q | q ↛ x ∈ E Σ(Γ,a) }
      If Γ* no longer contains a positive path from a to p,
        then remove the edge p • x from Σ(Γ, a)
    For each remaining positive edge p · x
      If Σ(Γ, a) contains a path p · q1 ··· qn → x, n ≥ 1
        such that there is no negative path from a to any of qi, x,
        then remove the edge p · x from Σ(Γ, a)
```

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

Acknowledgments

Discussions with Leora Morgenstern, David Touretzky, Jeff Horty, 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 *AAA],* Pages 104-108, Washingon, DC, 1983. Morgan Kaufmann Publishers, Inc.

[Horty *et ai,* 1987] John F. llorty, Richmond 11. 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, HI, 1989.

[Stein, 1990] Lynn Andrea Stein. Credulous extensions as possible worlds. *In Formal Aspects of Semantic Networks.* Morgan Kaufmann Publishers, Inc., 1990.

[Touretzky *et a/.,* 1987] David S. Touretzky, John F. Horty, and Richmond II. 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.

Knowledge Representation

Uncertainty, Belief, and Probability

Ronald Fagin and Joseph Y. Halpern
IBM Almaden Research Center, Dept. K53/S02
650 Harry Rd.
San Jose, CA 95120
email: fagin@ibm.com, halpern@ibm.com

Abstract

We introduce a new probabilistic approach to dealing with uncertainty, based on the observation that probability theory does not require that every event be assigned a probability. For a *nonmeasurable* event (one to which we do not assign a probability), we can talk about only the *inner measure* and *outer measure* of the event. Thus, the measure of belief in an event can be represented by an interval (defined by the inner and outer measure), rather than by a single number. Further, this approach allows us to assign a belief (inner measure) to an event E without committing to a belief about its negation $\neg E$ (since the inner measure of an event plus the inner measure of its negation is not necessarily one). Interestingly enough, inner measures induced by probability measures turn out to correspond in a precise sense to Dempster-Shafer belief functions. Hence, in addition to providing promising new conceptual tools for dealing with uncertainty, our approach shows that a key part of the important Dempster-Shafer theory of evidence is firmly rooted in classical probability theory.

1 Introduction

Dealing with uncertainty is a fundamental issue for AI. The most widely-used approach to dealing with uncertainty is undoubtedly the Bayesian approach. It has the advantage of relying on well-understood techniques from probability theory, as well as some philosophical justification on the grounds that a "rational'* agent must assign uncertainties to events in a way that satisfies the axioms of probability [Cox46, Sav54]. On the other hand, the Bayesian approach has been widely criticized for requiring an agent to assign a subjective probability to every event. While this can be done in principle by having the agent play a suitable betting game [Jef83],[1] it does have a number of drawbacks. Among others, there is the computational difficulty of arriving at the probability. There is also the issue of whether it is reasonable to describe confidence by a single point rather than a range. While an agent might be prepared to agree that the probability of an event lies within a given range, say between 1/3 and 1/2, he might not be prepared to say that it is precisely .435.

Not surprisingly, there has been a great deal of debate regarding the Bayesian approach (see [Che85] and [Sha76] for some of the arguments). Numerous other approaches to dealing with uncertainty have been proposed, including Demp>ster-Shafer theory [Dem68, Sha76], Cohen's model of endorsements [Coh85], and various nonstandard, modal, and fuzzy logics (for example, [HR87, Zad75]). A recent overview of the field can be found in [Saf88]. Of particular interest to us here is the Dempster-Shafer approach, which uses *belief {unctions* and *plausibility functions* to attach numerical lower and upper bounds on the likelihoods of events.

Although the Bayesian approach requires an agent to assign a probability to every event, probability theory does not. The usual reason that mathematicians deal with nonmeasurable events (those that are not assigned a probability) is out of mathematical necessity. For example, it is well known that if the sample space of the probability space consists of all numbers in the real interval [0, 1], then we cannot allow every set to be measurable if (like Lebesgue measure) the measure is to be translation-invariant (see [Boy64, page 54]). However, in this paper we allow nonmeasurable events out of choice, rather than out of mathematical necessity. An event E for which an agent has insufficient information to assign a probability is modelled as a nonmeasurable set. The agent is not forced to assign a probability to E in our approach. We can provide meaningful lower and upper bounds on our degree of belief in E by using the standard mathematical notions of *inner measun* and *outer measure* induced by the probability measure [Hal50], which, roughly speaking, are the probability of the largest measurable event contained in E and the smallest measurable event containing E, respectively.

Allowing nonmeasurable events has its advantages. The uncertainty of event E is no longer given by a single number, but rather by an interval defined by the inner and outer measures. Furthermore, it is possible for the belief (i.e., inner measure) of event E to be a without the belief of -^E being 1 — *a*. Rather than nonnieasurability being a mathematical nuisance, we have turned it

'This idea is due to Ramsey [Ram3I] and was rediscovered by von Neumann and Morgenstern [vNM47]; a clear exposition can be found in [LR57].

here into a desirable feature!

We feel that this paper makes three major contributions. The first is conceptual: In certain situations, our approach gives a useful way to think about and reason about uncertainty. In particular, the use of nonmeasurable sets seems to provide a useful way to capture our uncertainty about the probability of an event. The second is technical: We prove that, in a precise sense, inner measures induced by probability measures are equivalent to Shafer's belief functions (and so outer measures induced by probability measures are equivalent to Shafer's plausibility functions). The implications of this equivalence are significant. Although some, such as Cheeseman [Che85], consider the theory of belief functions as *ad hoc* and essentially nonprobabilistic (see discussion by Shafer [Sha86]), our results help show that a key part of the Dempster-Shafer theory of evidence is firmly rooted in classical probability theory. The last contribution is also technical: by combining our results here with those of a companion paper [FHM88], we are able to obtain a sound and complete axiomatization for a rich propositional logic of evidence, and provide a decision procedure for the satisfiability problem, which we show is no harder than that of propositional logic (NP-complete). Our techniques may provide a means for automatically deducing the consequences of a body of evidence.

2 Probability theory

To make our discussion precise, it is helpful to recall some basic definitions from probability theory (see [Fel57] for more details). A *probability space* $(S, \mathcal{X}, \mu)$ consists of a set S (called the *sample space),* a σ-algebra $\mathcal{X}$ of subsets of S (i.e., a set of subsets of S containing S and closed under complementation and countable union, but not necessarily consisting of all subsets of S) whose elements are called *measurable sets,* and a *probability measure* $\mu: \mathcal{X} \to [0, 1]$ satisfying the following properties:

P1. $\mu(X) \geq 0$ for all $X \in \mathcal{X}$

P2. $\mu(S) = 1$

P3. $\mu(\cup_{i=1}^{\infty} X_i) = \Sigma_{i=1}^{\infty} \mu(X_i)$, if the X_i's are pairwise disjoint members of $\mathcal{X}$.

Property P3 is called *countable additivity.* Of course, the fact that $\mathcal{X}$ is closed under countable union guarantees that if each $X_i \in \mathcal{X}$, then so is $\cup_{i=1}^{\infty} X_i$. If $\mathcal{X}$ is a finite set, then we can simplify property P3 above to

P3′. $\mu(X \cup Y) = \mu(X) + \mu(Y)$, if X and Y are disjoint members of $\mathcal{X}$.

This property is called *finite additivity.* Properties P1, P2, and P3′ characterize probability measures in finite spaces.

In a probability space $(S, \mathcal{X}, \mu)$, the probability measure μ is not defined on 2^S (the set of all subset of S), but only on $\mathcal{X}$. We can extend μ to 2^S in two standard ways, by defining functions μ_* and μ^*, traditionally called the *inner measure* and *outer measure induced by* μ [Hal50]. For an arbitrary subset $A \subseteq S$, we define

$$\mu_*(A) = \sup\{\mu(X) \mid X \subseteq A \text{ and } X \in \mathcal{X}\}$$
$$\mu^*(A) = \inf\{\mu(X) \mid X \supseteq A \text{ and } X \in \mathcal{X}\}.$$

If there are only finitely many measurable sets (in particular, if S is finite), then it is easy to see that the inner measure of A is the measure of the largest measurable set contained in A, while the outer measure of A is the measure of the smallest measurable set containing A. In any case, it is not hard to show by countable additivity that for each set A, there are measurable sets B and C where $B \subseteq A \subseteq C$ such that $\mu(B) = \mu_*(A)$ and $\mu(C) = \mu^*(A)$. Note that if there are no nonempty measurable sets contained in A, then $\mu_*(A) = 0$, and if there are no measurable sets containing A other than the whole space S, then $\mu^*(A) = 1$. The properties of probability spaces guarantee that if X is a measurable set, then $\mu_*(X) = \mu^*(X) = \mu(X)$. In general we have $\mu^*(A) = 1 - \mu_*(\overline{A})$.

Suppose we have a situation we want to reason about. Typically we do so by fixing a finite set $\Phi = \{p_1, \ldots, p_n\}$ of primitive propositions, which can be thought of as corresponding to basic events, such as "it is raining now" or "the coin landed heads". The set $\mathcal{L}(\Phi)$ of *(propositional) formulas* is the closure of Φ under the Boolean operations $\wedge$ and $\neg$. The primitive propositions in Φ do not in general describe mutually exclusive events. To get mutually exclusive events, we can consider all the *atoms*, that is, all the formulas of the form $p_1' \wedge \ldots \wedge p_n'$, where p_i' is either p_i or $\neg p_i$. Let At denote the set of atoms.

We have been using the word "event" informally, sometimes meaning "set" and sometimes meaning "formula". We now want to be more formal, and to be able to talk explicitly about the probability of a *formula*. However, a probability measure is a function on sets, not formulas. Fortunately, it is easy to go from sets to formulas.

Using standard propositional reasoning, it is easy to see that any formula can be written as a disjunction of atoms. Thus, a formula φ can be identified with the unique set $\{\delta_1, \ldots, \delta_k\}$ of atoms such that $\varphi \equiv \delta_1 \vee \ldots \vee \delta_k$. If we want to assign probabilities to all formulas, we can simply assign probabilities to each of the atoms, and then use the finite additivity property of probability measures to compute the probability of an arbitrary formula. This amounts to taking a probability space of the form $(At, 2^{At}, \mu)$. The states in the probability space are just the atoms, and the measurable subsets are all the sets of atoms (i.e., all formulas). Once we assign a measure to the singleton sets (i.e., to the atoms), we can extend by additivity to any subset. We call such a probability space a *Nilsson structure*, since this is essentially what Nilsson used to give meaning to formulas in his probability logic [Nil86].[2] Given a Nilsson structure $N = (At, 2^{At}, \mu)$ and a formula φ, let $W_N(\varphi)$ denote the *weight* or *probability* of φ in N, which is defined to be $\mu(At(\varphi))$, where $At(\varphi)$ is the set of atoms whose disjunction is equivalent to φ.

A more general approach is to take a *probability structure* to be a tuple $(S, \mathcal{X}, \mu, \pi)$, where $(S, \mathcal{X}, \mu)$ is a probability space, and π associates with each $s \in S$ a truth assignment $\pi(s): \Phi \to \{\mathbf{true}, \mathbf{false}\}$. We say that p is

[2] Actually, the use of possible worlds in giving semantics to probability formulas goes back to Carnap [Car50].

true at s if $\pi(s)(p)$ = true; otherwise, we say that *p* is *false at s.*

We think of *S* as consisting of the possible states of the world. We can associate with each state *s* in 5 a unique atom describing the truth values of the primitive propositions in *s.* For example, if Φ = *{p1,p2),* and if $\pi(s)(p_1)$ = true and $\pi(s)(p_2)$ = false, then we associate with *s* the atom $p_1 \wedge \neg p_2$. It is perfectly all right for there to be several states associated with the same atom (indeed, there may be an infinite number, since we allow *S* to be infinite, even though Φ is finite). This situation may occur if a state is not completely characterized by the events that are true there. This is the case, for example, if there are features of worlds that are not captured by the primitive propositions.

We can easily extend $\pi(s)$ to a truth assignment on all formulas by taking the usual rules of propositional logic. Then if *M* is a probability structure, we can associate with every formula φ the set φ^M consisting of all the states in *M* where φ is true (i.e., the set $\{s \in S \mid \pi(s)(\varphi) = \mathbf{true}\}$). Of course, we assume that *n* is defined so that $true^M$ — *S.* If p^M is measurable for every primitive proposition $p \in \Phi$, then φ^M is also measurable for every for inula φ (since the set *X* of measurable sets is closed under complementation and countable union). We say *M* is a *measurable probability structure* if φ^M is measurable for every formula φ.

It makes sense to talk about the probability of <p in *M* only if φ^M is measurable; we can then take the probability of φ, which we denote $W_M(\varphi)$, to be $\mu(\varphi^M)$. If $(f^{Xi}$ is not measurable, then we cannot talk about its probability. However, we can still talk about its inner measure and outer measure, since these are defined for all subsets. Intuitively, the inner and outer measure provide lower and upper bounds on the probability of φ. In general, if φ^M is not measurable, then we take $W_M(\varphi)$ to be $\mu_*(\varphi)$, i.e., the inner measure of φ in *M.*

We define a probability structure *M* and a Nilsson structure *N* to be *equivalent* if $W_M(\varphi)$ = $W_N(\varphi)$ for every formula φ. Intuitively, a probability structure and a Nilsson structure are equivalent if they assign the same probability to every formula. The next theorem shows that there is a natural correspondence between Nilsson structures and measurable probability structures.[3]

Theorem 2.1:

/. For every Nilsson structure there is an equivalent measurable probability structure.

2. For every measurable probability structure there is an equivalent Nilsson structure.

Why should we even allow nonmeasurable sets? As the following example shows (as do others given in the full paper), using nonmeasurability allows us to avoid assigning probabilities to those events for which we have insufficient information to assign a probability.

Example 2.2: Ron has two blue suits and two gray suits. He has a very simple method for deciding what color suit to wear on any particular day: he simply tosses a (fair) coin: if it lands heads he wears a blue suit, and if it lands tails he wears a gray suit. Once he's decided what color suit to wear, he just chooses the rightmost suit of that color on the rack. Both of Ron's blue suits are single-breasted, while one of Ron's gray suits is single-breasted and the other is double-breasted. Ron's wife Susan is (fortunately for Ron) a little more fashion-conscious than he is. She also knows how Ron makes his sartorial choices. So, from time to time, she makes sure that the gray suit she considers preferable is to the right (which it is depends on current fashions and perhaps on other whims of Susan).[4] Suppose we don't know about the current fashions (or about Susan's current whims). What can we say about the probability of Ron's wearing a single-breasted suit on Monday?

In terms of possible worlds, it is clear that there are four possible worlds, one corresponding to each of the suits that Ron could choose. For definiteness, suppose states $i and *s<>* correspond to the two blue suits, .s_3 corresponds to the single-breasted gray suit, and S4 corresponds to the double-breasted gray suit. Let *S* = {.s1, s2, S3, S4}. There are two features of interest about a suit: its color and whether it is single-breasted or double-breasted. Let the primitive proposition *g* denote "the suit is gray" and let *db* denote "the suit is double-breasted", and define the truth assignment *n* in the obvious way. Note that the atom ¬*g* /\ ¬*db* is associated with both states *s1* and *s2.* Since the two blue suits are both single-breasted, these two states cannot be distinguished by the formulas in our language.

What are the measurable events? Besides *S* itself and the empty set, the only other candidates are *{s_1,s_2}* ("Ron chooses a blue suit") and {.S3, s4} ("Ron chooses a gray suit"). However, *SB* — {s1, .s_2, .S3} ("Ron chooses a single-breasted suit") is nonmeasurable. The reason is that we do not have a probability on the event "Ron chooses a single-breasted suit, given that Ron chooses a gray suit", since this in turn depends on the probability that Susan put the single-breasted suit to the right of the other gray suit, which we do not know. Susan's choice might be characterizable by a probability distribution; it might also be deterministic, based on some complex algorithm which even she might not be able to describe; or it might be completely nondeterministic, in which case it is not technically meaningful to talk about the "probability" of Susan's actions! Our ignorance here is captured by nonmeasurability. Informally, we can say that the probability of Ron choosing a single-breasted suit lies somewhere in the interval [1/2,1], since it is bounded below by the probability of Ron choosing a blue suit. This is an informal statement because formally it does not make sense to talk about the probability of a nonmeasurable event. The formal analogue is simply that the inner measure of SB is 1/2, while its outer measure is 1. I

The proof of this and all other theorems mentioned here can be found in the full paper [FH88],

[4]Anv similarity between the characters in this example and the first author of this paper and his wife Susan is not totally accidental.

3 The Dempster-Shafer theory of evidence

The Dernpster-Shafer theory of evidence [Sha76] provides another approach to attaching likelihoods to events. This theory starts out with a *belief function* (sometimes called a *support function).* For every event (i.e., set) *A,* the belief in *A,* denoted *Bel(A),* is a number in the interval [0,1] that places a lower bound on likelihood of *A.* We have a corresponding number *Pl(A)* = 1 — *Bel(A),* called the *plausibility* of *A,* which places an upper bound on the likelihood of *A.* Thus, to every event *A* we can attach the interval *[Bel(A), P1(A)].* Like a probability measure, a belief function assigns a "weight" to subsets of a set 5, but unlike a probability measure, the domain of a belief function is always taken to be *all* subsets of *S.* Just as we defined probability structures, we can define a *DS structure* (where, of course, "DS" stands for Dempster-Shafer) to be a tuple *(S, Bel,* TT), where *S* and are as before, and where *Bei.*2^s —> [0, 1] is a function satisfying:

B1. $Bel(\emptyset) = 0$

B2. $Bel(S) = 1$

B3. $Bel(A_1 \cup \ldots \cup A_k) \geq \sum_{I \subseteq \{1,\ldots,k\}, I \neq \emptyset} (-1)^{|I|+1} Bel(\bigcap_{i \in I} A_i)$.

A belief function is typically defined on a *frame of discernment,* consisting of mutually exclusive and exhaustive propositions describing the domain of interest. We think of the set *S* of states in a belief structure as being this frame of discernment. We could always choose *S* to be some subset of *At,* the set of atoms, so that its elements are in fact propositions in the language. In general, given a DS structure $D = (S, Bel, \pi)$ and formula φ, we define the weight $W_D(\varphi)$ to be $Bel(\varphi^D)$, where φ^D is the set of states where *(p* is true. Thus we can talk about an agent's degree of belief in <p in *D,* described by $W_D(\varphi)$, by identifying <p with the set φ^D and considering the belief in φ^D. As before, we define a probability structure *M* (resp., a Nilsson structure *N,* a DS structure D') and a DS structure D to be *equivalent if* $W_M(\varphi) = W_D(\varphi)$ (resp., $W_N(\varphi) = W_D(\varphi)$, $W_{D'}(\varphi) = W_D(\varphi)$) for every formula φ.

Property B3 may seem unmotivated. Perhaps the best way to understand it is as an analogue to the usual inclusion-exclusion rule for probabilities [Fel57, p. 89], which is obtained by replacing the inequality by equality (and the belief function *Bel* by a probability measure). In particular, B3 holds for probability measures (we prove a more general result, namely that it holds for all inner measures induced by probability measures, in Proposition 3.1 below). Hence, if *(S, X, µ)* is a probability space and *X* = 2^s (making every subset of 5 measurable), then µ is a belief function. (This fact has been observed frequently before; see, for example, [Sha76].) It follows that every Nilsson structure is a DS structure.

It is easy to see that the converse does not hold. For example, suppose there is only one primitive proposition, say p, in the language, so that *At* = {p, -p}, and let D_0 = *(At, Bel,* π) be such that *Bel({p})* = 1/2, $Bel(\{\neg p\})$ = 0, and π is defined in the obvious way. Intuitively, there is weight of evidence 1/2 for p, and no evidence for -p. Thus $W_{D_0}(p)$ = 1/2 and $W_{D_0}(\neg p)$ = 0. D_0 is not equivalent to any Nilsson structure, since if *N* is a Nilsson structure such that *WN(P)* = 1/2, then we must have $W_N(\neg p)$ = 1/2.

These observations tell us that in some sense belief functions are more general than probability measures, provided we restrict attention to probability spaces where all sets are measurable. This fact is well known. Indeed, in [Sha76], Shafer makes explicit use of the greater generality of belief functions. While he does consider events *E* such that *Bel(¬E)* = 1 - *Bel(E)* (he calls such events *probabilistic),* he also wants to allow non-probabilistic events. He gives examples of events where the fact that we would like to assign weight .8 to our belief in event *E* does not mean that we want to assign weight .2 to our belief in *Ì¬E.* In our framework, where we allow nonmeasurable sets, we can view probabilistic events as corresponding to measurable sets, while non-probabilistic events do not. We can push this analogy much further. Not only do nonmeasurable sets correspond to non-probabilistic events, but the inner measures induced by probability measures correspond to belief functions.

Proposition 3.1: *// (5, X, p) is a probability space, then* μ_* *is a belief function on* 2^S .

Proposition 3.1 says that every inner measure is a belief function (and thus generalizes the statement that every probability measure is a belief function). The converse does not quite hold. For example, consider the DS space D_0 defined above. There is no probability measure // that we can define on {p, ~~>p) such that $\mu_* = Bel$. However, it is easy to define a probability structure *M* such that $\mu_*(p^M) = 1/2$ and $\mu_*(\neg p^M) = 0$. That is, we can find a probability structure equivalent to Do- The next theorem generalizes this observation.

Theorem 3.2:

1. *For every DS structure there is an equivalent probability structure.*
2. *For every probability structure there is an equivalent DS structure.*

Intuitively, Theorem 3.2 says that belief functions and inner measures induced by probability measures are precisely the same if their domains are considered to be formulas rather than sets. As we shall see, this result has important implications regarding complete axiomatizations and decision procedures.

4 Reasoning about belief and probability

We are often interested in the inferences we can make about probabilities or beliefs given some information. In order to do this, we need a language for doing such reasoning. Such a language is given in [FHM88]. A *term* in this language is an expression of the form $a_1 w(\varphi_1) + \cdots + a_k w(\varphi_k)$, where $a_1, \ldots, a_k$ are integers and $\varphi_1, \ldots, \varphi_k$ are propositional formulas. A *basic weight formula* is one of the form $t \geq b$, where *t* is a term and b is an integer. A *weight formula* is a Boolean combination of basic weight formulas. We sometimes use obvious abbreviations without further comment, such as $w(\varphi) \geq w(\psi)$ for $w(\varphi) - w(\psi) \geq 0$.

We give semantics to the formulas in our language with respect to all the structures we have been considering. Let A' be either a Nilsson structure, a probability structure, or a DS structure, and let / be a weight formula. We now define what it means for K to *satisfy* /, written $K \models f$. For a basic weight formula,

$$K \models a_1 w(\varphi_1) + \cdots + a_k w(\varphi_k) \geq b \text{ iff } \\ a_1 W_K(\varphi_1) + \cdots + a_k W_K(\varphi_k) \geq b.$$

We then extend $\models$ in the obvious way to conjunctions and negations. The interpretation of $w(\varphi)$ is either "the probability of φ" (for Nilsson structures or measurable probability structures), "the inner measure o φ" f o r general probability structures), or "the belief in φ" (for DS structures).

Let K be a class of structures (in the cases of interest to us, K is the class of either probability structures, measurable probability structures, Nilsson structures, or DS structures). As usual, we define a weight formula / to be *satisfialet with nsptct to* $\mathcal{K}$ if $K \models f$ for some $K \in \mathcal{K}$. Similarly, / is *valid with respect to* $\mathcal{K}$ if $K \models f$ for all $K \in \mathcal{K}$

In [FHM88], an axiom system *AX ME AS FOR* reasoning about measurable probability structures is provided. The system has three parts, which deal respectively with propositional reasoning, reasoning about linear inequalities, and reasoning about probability. For example, a typical axiom for reasoning about linear inequalities is

$$(a_1 w(\varphi_1) + \cdots + a_k w(\varphi_k) \geq b) \Rightarrow \\ (ca_1 w(\varphi_1) + \cdots + ca_k w(\varphi_k) \geq cb) \text{ if } c \geq 0,$$

which says that both sides of an inequality can be multiplied by a positive constant. (The remaining axioms for reasoning about inequalities are described in the full paper.)

For reasoning about probability, we have the following axioms. The first three correspond to the usual laws of probability, except that W3 corresponds to finite additivity, not countable additivity.

W1. $w(\varphi) \geq 0$ (nonnegativity)

W2. $w(true) = 1$ (the probability of the event $true$ is 1)

W3. $w(\varphi \wedge \psi) + w(\varphi \wedge \neg\psi) = w(\varphi)$ (additivity)

W4. $w(\varphi) = w(\psi)$ if $\varphi \equiv \psi$ is a propositional tautology

As is shown in [FHM88], AXMEAS characterizes the valid formulas for measurable probability structures.

Theorem 4.1: ([FHM88]) AX_{MEAS} *is a sound and complete axiomatization for weight formulas with respect to measurable probability structures.*

This result, together with Theorem 2.1, immediately gives us

Corollary 4.2: *AXMEAS IS A sound and complete arlomatizatwn for weight formulas with respect to Nilsson structures.*

AX MEAS is not sound with respect to arbitrary probability structures, where $w(\varphi)$ is interpreted as the inner measure of φ^M. In particular, axiom W3 no longer holds: inner measures are not finitely additive. Let AX be obtained from AXMEAS by replacing W3 by the following two axioms, which are obtained from conditions BI and B3 for belief functions in an obvious way:

W5. $w(false) = 0$

W6. $w(\varphi_1 \vee \ldots \vee \varphi_k) \geq \\ \sum_{I \subseteq \{1,\ldots,k\}, I \neq \emptyset} (-1)^{|I|+1} w(\bigwedge_{i \in I} \varphi_i)$

Theorem 4.3: ([FHM88]) *AX is a sound and complete axiomatization for weight formulas with respect to probability structures.*

Applying Theorem 3.2, we immediately get

Corollary 4.4: *AX is a sound and complete automatization for weight formulas with respect to DS structures.*

Thus, using AX, we can derive all consequences of a collection of beliefs.

Combining the preceding results with results of [FHM88], we can also characterize the complexity of reasoning about probability and belief.

Theorem 4.5: *The complexity of deciding whether a weight formula is satisfiable with respect to probability structures (respectively, measurable probability structures, Nilsson structures, DS structures) is NP-complete.*

(This result in the ca.se of Nilsson structures was obtained independently in [GKP88].) Note that Theorem 4.5 says that reasoning about probability and belief is, in a precise sense, exactly as difficult as propositional reasoning. This is the best we could expect, since it is easy to see that reasoning about probability and belief is at least as hard as propositional reasoning (the propositional formula φ is satisfiable iff the weight formula $w(\varphi) > 0$ is satisfiable).

5 Combining evidence

An important issue for belief functions, each of which can be viewed as representing a distinct body of evidence, is how to combine them to obtain a new belief function that somehow reflects the combined evidence. A way of doing so is provided by Dempster's *rule of combination,* which was introduced by Dempster [Dem68] and was further developed and studied in an elegant and rather complete manner by Shafer [Sha76].

In the full paper [FH88], we show that there is a natural way (in the spirit of Dempster's rule) to define the combination $D_x \oplus D_2$ of two DS structures D_1 and D_2, and a natural way to define the combination M_1 ® M_2 of two probability structures M_1 and M_2, such that the following theorem holds.

Theorem 5.1: *Let D_1 and D_2 bt DS structures. There are probability structures M_1 and M_2 such that (a) D\ is equivalent to M_1, (b) D_2 is equivalent to M_2, and (c) D\ $\oplus$ D'2 is equivalent to M\ $\oplus$ A/$_2$.*

This theorem shows that the spirit of Dempster's rule of combination can be captured within probability theory We are currently investigating alternatives to Dempster's rule of combination for revising beliefs about uncertainty in the presence of new information. The idea is to consider what it means to take a conditional probability with respect to a nonmeasurable set. We plan to report on this work in a future paper.

6 Related work

Although we believe we are the first to propose using inner and outer measures as a way of dealing with uncertainty, there are a number of other works with similar themes. We briefly discuss them here.

A number of authors have argued that we should think in terms of an interval in which the probability lies, rather than a unique numerical probability (see, for example, [Kyb61, Kyb88]). Good [Goo62], Koopman [Koo40a, Koo40b], and Smith [Smi61] try to derive reasonable properties for the intuitive notions of *lower* and *upper* probability, which are somehow meant to capture lower and upper bounds on an agent's belief in a proposition. Good observes that "The analogy [between lower and upper probability and] inner and outer measure is obvious. But the axioms for upper and lower probability do not all follow[7] from the theory of inner and outer measure."

Dempster [Dem66, Dem68] gives a formal mathematical definition of lower and upper probability in terms of a tuple (,S,μ,T,T), which we call a *Dempsler strvctare. (S,* 2[s] ,μ is a probability space (Dempster assumes for simplicity that every subset of *S* is measurable). *T* is another set, and F is a multi-valued mapping *S* to *T.* Thus, *T(s)* is a subset of *T* for each *s* *S.* Given $A \subseteq T$, we define subsets A, and *A** of *S* as follows:

$$A_* = \{s \in S \mid \Gamma(s) \neq \emptyset, \Gamma(s) \subseteq A\}$$
$$A^* = \{s \in S \mid \Gamma(s) \cap A \neq \emptyset\}.$$

Provided $\mu(T^*) \neq 0$, we define the lower and upper probabilities of A, written $P_*(A)$ and $P^*(A)$ respectively, by

$$P_*(A) = \mu(A_*)/\mu(T^*)$$
$$P^*(A) = \mu(A^*)/\mu(T^*).$$

It is easy to check that $T_* = T^*$. Thus, dividing by $\mu(T^*)$ has the effect of normalizing so that $P_*(T) = P^*(T) = 1$.

There is a close relationship between lower and upper probabilities and inner and outer measures induced by a probability measure. Given a probability structure $M = (S, \mathcal{X}, \mu, \pi)$ where S is finite, let $(\mathcal{X}', \mu', T, \Gamma)$ be the Dempster structure where (1) $\mathcal{X}'$ is a basis for $\mathcal{X}$,[5] (2) μ' is a probability measure defined on $2^{\mathcal{X}'}$ by taking $\mu'(\{A\}) = \mu(A)$ for $A \in \mathcal{X}'$ and then extending to all subsets of $\mathcal{X}'$ by finite additivity, (3) T consists of all propositional formulas, and (4) for $A \in \mathcal{X}'$, we define F(A) to consist of all formulas φ such that φ is true at some point in .4 (in the structure A/). Thus T is a multivalued mapping from *X'* to *T.* It is easy to check that $P_*(\{\varphi\}) = \mu_*(\varphi^M)$ and $P^*(\{\varphi\}) = \mu^*(\varphi^M)$ for all formulas φ

Ruspini [R.us87] also considers giving semantics to probability formulas by using possible worlds, but he includes epistemic notions in the picture. Briefly, his approach can be described as follows (where have taken the liberty of converting some of his notation to ours, to make the ideas easier to compare). Fix a set *{p₁,.. ., pₙ]* of primitive propositions. Instead of considering just propositional formulas, Ruspini allows epistemic formulas; he obtains his language by closing off under the propositional connectives ∧, ∨, ⇒, and , ¬ as well as the epistemic operator A". Thus, a typical formula in his language would be $K(p_1 \Rightarrow K(p_2 \wedge p_3))$. (A formula such as $K\varphi$ should be read "the agent knows) Rather than considering arbitrary sample spaces as we have done here, where at each point in the sample space some subset of primitive propositions is true, Ruspini considers one fixed sample space *S* (which he calls a *sentence space)* whose points consist of all the possible truth assignments to these formulas consistent with the axioms of the modal logic S5. (See, for example, [HM85] for an introduction to S5. We remark that it can be shown that there are less than 2[r)]2[2] consistent truth assignments, so that 5 is finite.) We can define an equivalence relation — on *S* by taking .s ~ *t* if *s* and */* agree on the truth values of all formulas of the form $K\varphi$. The equivalence classes form a basis for a σ-algebra of measurable subsets of *S.* Let *X* be this σ-algebra. For any formula φ, l e φ^S o n s i s t of all the truth assignments in *S* that make φ true. It is easy to check that $(K\varphi)^S$, the set of truth assignments that make $K\varphi$ true, is the union of equivalence classes, and hence is measurable. Let // be any probability measure defined on *X.* Given /i, we can consider the probability structure $(S, \mathcal{X}, \mu, \pi)$, where we take $\pi(s)(p) = s(p)$. (Since .s is a truth assignment, this is well defined.) The axioms of S5 guarantee us that *{F\<p)[s]* is the largest measurable subset contained in φ^M; thus, $\mu_*(\varphi^M) = \mu((K\varphi)^S)$.

Ruspini then considers the DS structure (At, Bel, π'), where π' is defined in the obvious way on the atoms in ,4/, and $Bel(\varphi^D) = \mu((K\varphi)^S)(= \mu_*(\varphi^M))$.[6] Ruspini shows that *Bel* defined in this way is indeed a belief function. Thus, Ruspini shows a close connection between probabilities, inner measures, and belief functions in the particular structures that he considers. He does not show a general relationship between inner measures and belief functions; in particular, he does not show that DS structures are equivalent to probability structures, as we do in Theorem 3.2.

In the full paper, we explore further relations between our work and that of Ruspini, as well as comparing our characterization of belief functions with those of Shafer [Sha79], Kyburg [Kyb87], and Pearl [Pea88].

'A subset *X'* of *X* is said to be a *basis* (of *X)* if the members of *X'* are nonempty and disjoint, and if *X* consists precisely of countable unions of members of *X'.* It is easy to see that if *X* is finite then it has a basis. Moreover, whenever *X* has a basis, it is unique: it consists precisely of the minimal elements of *X* (the nonempty sets none of whose nonempty subsets are in *X).* Note that if *X* has a basis, once we know the probability of every set in the basis, we can compute the probability of every measurable set by using countable additivity.

[6]Ruspini actually defines the belief function directly on formulas; i.e., he defines $Bel(\varphi)$. In our notation, what he is doing is defining a weight function *WD-*