An Evidence-based Framework for a Theory of Inheritance

T. Krishnaprasad, Michael Kifer
Department of Computer Science
SUNY at Stony Brook
Stony Brook, NY 11794.

Abstract

We present an approach to formalizing non-monotonic multiple inheritance networks by combining concepts from logic programming and multi-valued logics in a uniform framework. A Horn-clause logic language is used for specifying inheritance networks. This allows a natural representation of class-subclass hierarchies and ambiguous inheritance networks. It also provides means for resolving ambiguities resulting from the network topology, but which are not inherent to the problem. We provide a model theory for the language and show how a unique intended model can be associated with every inheritance network. This model resembles the unique extension obtained in the skeptical theory of inheritance [Hor-87], but is more general. Finally, we present an algorithm which realizes the aforementioned semantics.

1 Introduction

The notion of nonmonotonic inheritance is fundamental to common-sense reasoning. For instance, knowing that *Bateman* is a mammal, one would conclude that it does not fly. This is because normally mammals do not fly and in the absence of any other information, we regard *Bateman* as inheriting its inability to fly from mammals. If one later learns that *Bateman* is a bat, then he will probably change his mind, concluding that it does fly. However, after learning that *Bateman* is a dead bat, he would again change his mind concluding that it cannot fly. In this example, knowing that *Bateman* is a bat is more informative than knowing that it is a mammal, and knowing that it is a dead bat is even more informative, as far as *Bateman's* ability to fly is concerned. In essence, the knowledge that an individual belongs to a subclass provides more information about the individual than the knowledge that the individual belongs to its superclass. We also notice that the first two conclusions are defeasible, while the third one is not.

Birds normally fly, while toys normally do not. If an item is a toy bird, then we conclude that it does not fly. This is because the item being a toy contributes more evidence in support of its inability to fly than does bird in support of its flying ability. Notice that there is no class-subclass relationship between toys and birds.

Some approaches to inheritance (e.g., [Touretzky, 1986], [Horty *et* a/., 1987]) are proof-theoretic. They give algorithms for computing sets of *acceptable* paths supported by a network, rather than specifying the states of the world the network represents. Others (e.g., [Etherington, 1983], [Haugh, 1988], [Krishnaprasad *et* a/., 1988a], [Przymusinska and Gelfond, 1988], [Th omason *ei* a/., 1987]) present translations of inheritance networks into some standard logical formalism. The semantics of the networks is then captured through the model theory for the respective logical formalism. In particular, a theory based on prioritized circumscription transforms a network into a set of first-order sentences augmented with meta-level minimality constraints embodying *preferences* [Krishnaprasad *et* a/., 1988a]. The first-order models of the resulting translation are the states of the world the network represents. The circumscriptive theory of [Haugh, 1988] formalizes the network at a meta-level. The set of models of the translation in [Haugh, 1988] *encodes* the meaning of the network. [Przymusinska and Gelfond, 1988] views a network as representing a set of beliefs of a rational agent, and captures this interpretation by translating the network into Moore's Autoepistemic logic [Moore, 1985]. The preference criteria is axiomatized in the translated theory. [Pearl, 1988] provides probabilistic semantics to networks, by assigning to them a set of possible worlds and an associated probability distribution.

We propose to view networks as specifying a set of belief-evidence pairs. "Strength of evidence" is explicitly incorporated into the object language, and we have developed a model theory for the resulting logic. The dominance of property inheritance from a subclass over that from a superclass is captured by making the evidence contributed by subclass membership stronger than that contributed by superclass membership. This provides an evidence based semantics to networks as a set of justified beliefs.

In this paper, we present a logic to formalize inheritance networks by combining concepts from logic programming and multi-valued logics in a uniform framework. The main ideas behind this approach are as follows:

- A naive formalization of inheritance networks in first-order logic leads to inconsistencies. As we do not have a satisfactory means of modifying a theory in the presence of "equally strong" conflicting evidences, we pursue an approach in which contradiction can be represented explicitly [Belnap, 1977a] [Blair, 1987] [Kifer and Lozinskii, 1989] [Thomason *et* a/., 1987], thereby providing an adequate representational framework for ambiguous inheritance networks.

- Ambiguity in inheritance networks leads to an explosion in the number of expansions that could arise. (See [Etherington, 1983] [Haugh, 1988] [Krishnaprasad *et* a/., 1988a] [Padgham, 1988] [Przymusinska and Gelfond, 1988] [Touretzky, 1986].) Furthermore, reasoning in the intersection of all the expansions is not computationally viable [Touretzky, 1987]. Thus, computational efficiency considerations force us to look in the direction of formalisms that associate fewer models to networks. In this paper, we show how a unique model can be associated with every inheritance network. This resembles the *skeptical* theory of non-monotonic inheritance networks of [Horty *et* a/., 1987], but is more general.
- To compute property inheritance, we need to perform certain bookkeeping functions such as tracking down the relative strengths of evidences. This is formalized by specifying how true or false a given statement is and how much is known about it [Ginsberg, 1986].
- Our specification language is a Horn clause language. The "ordering" of the evidences in the rule heads embodies the meta-knowledge about the priorities associated with the abnormality predicates in the circumscriptive theory of [Krishnaprasad *et* a/., 1988a]. The truth and the information orderings can be incorporated into the logic language by extending its syntax and modifying the definition of satisfaction, that is, truth of a sentence in a model, along the same lines as in [Blair, 1987] [Kifer and Li, 1988] [Kifer and Lozinskii, 1989].
- Most importantly, our approach is based on solid logic foundations, which provides a basis for the design and optimization of provably correct inheritance algorithms.

In summary, the evidence-based logic developed here gives an alternative way of representing inheritance networks. In addition, this formalism can also be used to specify the formal semantics of inheritance networks by translating the networks into it.

In Section 2, we develop a logical framework for inheritance networks. In Section 3, we discuss the characteristics of our approach and present extensions such as "preferential" inheritance and inheritance through paths with "negative" arcs.

2 An Evidence-Based Theory

2.1 Language

A *term* is an individual constant or a variable. An *atom* is a proposition or a formula q(t), where q is a unary predicate and t is a term. *Literals* are of the form p: r, where p is an atom and r is a *priority constant.* A *rule* is a statement of the form $p : r \Leftarrow q : \gamma$, where p and q are atoms, r and are priority constants. A *fact* is a ground literal. A *clause* is either a fact or a rule.

To capture dependency relationship among predicates, we define $\prec$ relation as follows.

Definition 1 q $\prec$ p (read q precedes p) iff either there is a rule in P of the form $p : r \Leftarrow q : \gamma$, or, recursively, $p : r \Leftarrow r : \lambda$ and q $\prec$ r.

An *inheritance specification* is a set of clauses, such that $\prec$ is an acyclic partial order on atoms. This eliminates the possibility of "recursive" inheritance, whose utility is unclear.

2.2 Priority constants

The priority constant r in p : r represents the type and the relative strength of evidence in support of p. Following [Belnap, 1977a] [Ginsberg, 1986J [Fitting, 1988], the priority constants are viewed on two different scales: one on the basis of their truth-content and the other on the basis of their information-content. For instance, —3 in the negative literal p : —3 represents an evidence of strength 3 supporting -p. The following relations are defined on the set of priority constants.

- <*k* is a semi-lattice order, i.e., a partial order equipped with the least upper bound operation. We will use *lubk(A)* to denote the least upper bound of the set *A.*
- ≈t is an equivalence relation on priority constants with respect to truth-content. For our purposes, we need only three equivalence classes representing differing information levels of truth, falsity or inconsistency in the evidence for an atom. To model both strict and defeasible links, we require that each equivalence class of $\approx_t$ has a unique $<_k$-maximal element. These can be thought of as signifying *true, false* or *top* (inconsistent) element and may be used to override any default conclusions that may be derivable from the specification.

Additionally, we define:

- $(\tau <_k \gamma)$ iff $(\tau \leq_k \gamma) \wedge (\tau \neq \gamma)$.
- $(\tau \leq_{tk} \gamma)$ iff $(\tau \leq_k \gamma) \wedge (\tau \approx_t \gamma)$.

2.3 Interpretation and Model

Let E be a set of clauses. The *domain* D of any Herbrand interpretation is collection of all individual constants mentioned in E. A *Herbrand base* of E is a collection of all facts of the form p : r, where p is a ground atom and r is a priority constant. A *Herbrand interpretation* I of E is a subset of the Herbrand base of E, such that if $p : \tau \in$ I and $\lambda \leq_{tk} \tau$ then $p : \lambda \in$ I. Furthermore, we require that if $p : \lambda \in$ I and $p : \tau \in$ I then $\lambda \approx_t \tau$.

Given two interpretations I and J, J $\sqsubseteq_k$ I (resp. J $\sqsubseteq_{tk}$ I) iff for every literal $p : \tau \in$ J, there exists a literal $p : \lambda \in$ I such that $\tau \leq_k \lambda$ (resp. $\tau \leq_{tk} \lambda$). Note that J $\sqsubseteq_{tk}$ I is equivalent to saying that J $\subseteq$ I. An interpretation I is *minimal* in a set of interpretations $\mathcal{S}$ iff $\neg\ \exists$J $\in \mathcal{S}$: J $\sqsubset_k$ I and it is *maximal* iff $\neg\ \exists$J $\in \mathcal{S}$: I $\sqsubset_k$ J.

A fact $p(t) : \tau$ is *satisfied* in I (denoted I $\models_k p(t) : \tau$) iff there is a fact $p(t) : \lambda \in$ I such that $\tau \leq_k \lambda$. A fact $p(t) : \tau$ is *strongly satisfied* in I (denoted I $\models_{tk} p(t) : \tau$) iff $\exists\ p(t) : \lambda \in$ I such that $\tau \leq_{tk} \lambda$. Equivalently, I $\models_{tk} p(t) : \tau$ iff $p(t) : \tau \in$ I.

A ground rule $p : \tau \Leftarrow q : \gamma$ is *satisfied* in I iff whenever I $\models_{tk} q : \gamma$, I $\models_k p : \tau$. A nonground rule $p : \tau \Leftarrow q : \gamma$ is *satisfied* in I iff all its ground instances arc satisfied in I. (A fact may be thought of as a rule whose empty body is strongly satisfied in all interpretations.)

A set of clauses P is *satisfied* in I, if all its clauses are satisfied in I. A literal $p : \lambda$ is *supported* by P in I if $\lambda \leq_{tk} lub_k(\{\gamma \mid p : \gamma \Leftarrow q : \delta \in$ and I $\models_{tk} q : \delta\})$. In words, p : A is supported by P in I if the combined conclusion derived from P on the basis of assuming facts in I strongly satisfies $p : \lambda$.

An interpretation I is a *model* of a set of clauses P iff P is satisfied under I. A model I of P is *supported* if every literal $p : \tau \in$ I is supported by P in I.

2.4 Existence of Supported Models

In this section, we show that every inheritance specification has supported models.

Lemma 1 *Every inheritance specification* P *consisting only of facts admits a supported Herbrand model.*

By *stratifying* P based on $\prec$-ordering, a supported model can be constructed in a standard way, using techniques similar to those described in [Apt *et al,* 1987].

Lemma 2 *Every inheritance specification* P *admits a supported Herbrand model.*

2.5 Unique model semantics

In this section, we demonstrate that a unique supported model can be naturally associated with every inheritance specification.

Lemma 3 *For an inheritance specification,* P, *consisting of facts only, there exists a unique supported Herbrand model of* P.

Given an inheritance specification P, we define an operator *Tp* that maps a Herbrand interpretation I into a Herbrand interpretation Tp(I) as follows.

Definition 2 **$\mathcal{T}_P(I)$ is a $\sqsubseteq_k$-minimal interpretation satisfying the set of literals $\{ p : \tau \mid p : \tau \Leftarrow q : \gamma$ is a ground instance of a clause in P and I $\models_{tk} q : \gamma \}$.**

Note that $\mathcal{T}_P$ is not monotonic with respect to $\sqsubset_k$. For instance, let P consist of a single rule $p : -1 \Leftarrow q : +1$ and consider a pair of interpretations I and J such that I $= \{p : -1, q : +2\}$ and J $= \{p : -2, q : -3\}$. (See Figure 1.) Clearly, I $\sqsubseteq_k$ J. However, $\mathcal{T}_P(I) = \{p : -1\}$ and $\mathcal{T}_P(J) = \emptyset$. Fortunately, $\mathcal{T}_P$ has the following weaker property that saves the situation: if I $\sqsubseteq_{tk}$ J then $\mathcal{T}_P(I) \sqsubseteq_k \mathcal{T}_P(J)$.

Lemma 4 *$\mathcal{T}_P(I) \sqsubseteq_{tk}$ I iff I is a model of* P, *and $\mathcal{T}_P(I)$ = I iff I is a supported model of* P.

We show the existence and uniqueness of a supported model of an inheritance specification P in two steps.

Theorem 1 *Every inheritance specification* P *admits at most one supported model Equivalently, there exists at most one solution to the fixpoint equation Tp(I) = I.*

The unique supported model for P, denoted M p , can be described $\varepsilon \lim_{i\to\infty} \mathcal{T}^i_P(\emptyset)$ ut this does not give u s an efficient way of computing the fixpoint of 7p because the successive approximations change nonmonotonically. However, we can build Mp monotonically by an iterated fixpoint construction, similar to [Apt *et a*/, 1987].

For convenience of exposition, and wig., we assume that predicates appearing as facts and those that appear as heads of the rules are disjoint.

<u>Inheritance Algorithm</u>

Let H be the set of ground atoms, A_i be the set of atoms whose literals acquire meaning at the i^{th} step of the iteration, and M_i be the set of literals at the i^{th} step of the iteration representing the intermediate state in the computation of the supported model for P. Then

- $\mathcal{A}_0 = \emptyset$.
- **$\mathcal{M}_0$ is the $\sqsubseteq_k$-minimal interpretation satisfying all ground facts in P.**
- **$\mathcal{A}_{i+1} = \mathcal{A}_i \cup \{p \mid p$ is a $\prec$-minimal ground atom in $(\mathbf{H} - \mathcal{A}_i)\}$. (Recall: $\prec$ is the "depends on" relation defined in Section 2.1.)**
- **$\mathcal{M}_{i+1}$ is the $\sqsubseteq_k$-minimal interpretation satisfying $\mathcal{M}_i \cup \{p : \tau \mid p$ is a $\prec$-minimal ground atom in $(\mathbf{H} - \mathcal{A}_i)$, $p : \tau \Leftarrow q : \gamma$ is a ground instance of a clause in P and $\mathcal{M}_i \models_{tk} q : \gamma\}$.**
- $\mathcal{M}_P = \bigcup_i \mathcal{M}_i$.

Theorem 2 *The set of literals Mp computed by the algorithm is a supported model of P.*

Note that absence of function symbols in our language makes our approach computationally viable.

From Theorem 1 and Theorem 2 it follows that

Theorem 3 *Every inheritance specification P has a unique supported model Mp associated with it.*

3 Discussion

In this section, we illustrate our theory on a number of examples. In particular, we explain how this language can be used to specify inheritance networks. For definiteness let us pick the semi-lattice in Figure 1. The priority constants are numbers annotated with symbols +, —, or * . All the constants with the same sign are truth-equivalent w.r.t. $\approx_t$ (i.e., —1 $\approx_t$ —3, *2 $\approx_t$ *5, etc.). Intuitively, constants +1, +2,... denote the degree of confidence that a fact is true, while —1, —2,... denote the degree of confidence in the falsehood of a fact. The constants *1, *2,... represent the degree of inconsistency in the evidences supporting facts. Also +w, —w, and *w represent *certain* information about truth, falsity and inconsistency respectively.

Consider the <k-ordering depicted in Figure 1. The meaning resulting from such a choice resembles the meaning attributed to networks by the skeptical theory

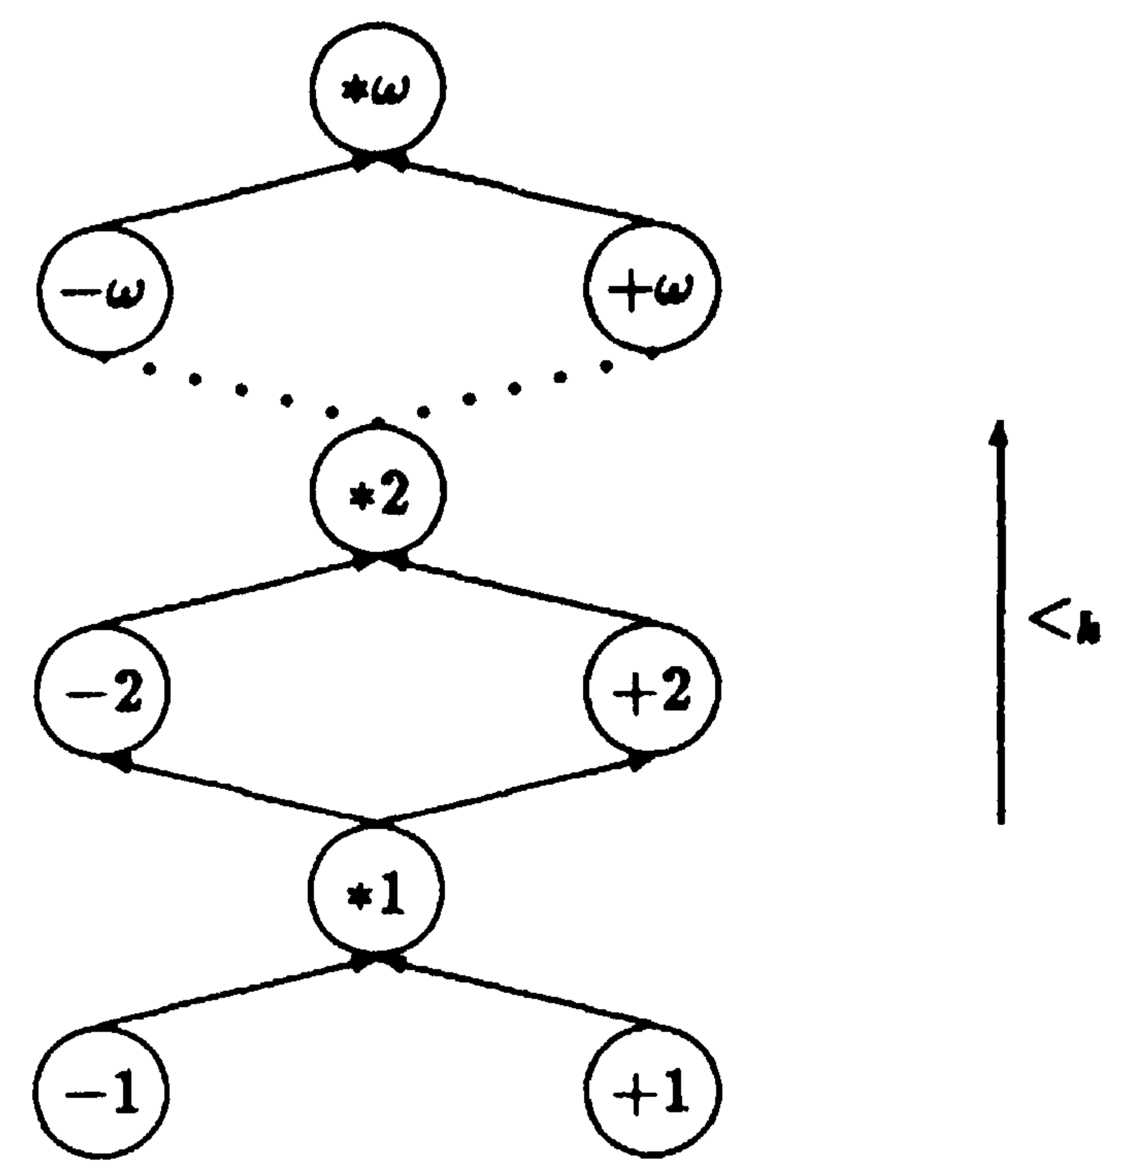

Figure 1: Information ordering for skeptical semantics

of [Horty *et al.*, 1987]. Intuitively, $p(a) : +n \in \mathcal{M}_P$ means that "*a* has property p", and $p(a) : -m \in \mathcal{M}_P$ stands for "*a* has property -p". In contrast, $p(a) : *l \in \mathcal{M}_P$ is interpreted as ambiguity about *p(a)*.

To evaluate a ground atomic query *p* with respect to an inheritance specification P, determine the maximal support, r, of a corresponding literal $p : r$ in $\mathcal{M}_P$. Note that this is well- defined by the definition of an interpretation.

Our inheritance networks are *bipolar* DAGs, that is, both positive and negative arcs are present. Our formalism is powerful enough to specify inheritances of individuals known to possess negated properties. Furthermore, we take the view of inheritance in which individuals are "moving upwards" similar to the one taken in [Horty *et al.*, 1987] and [Krishnaprasad *et al.*, 1988a]. The properties of an individual are obtained by forward chaining and there is no coupling among the inheritances of individuals belonging to the same class (unlike [Touretzky, 1987]).

Paraphrasing the famous Tweety example, we get the following rules.

- $fly(x) : +1 \Leftarrow bird(x) : +1.$
- $fly(x) : -2 \Leftarrow penguin(x) : +1.$
- $bird(x) : +\omega \Leftarrow penguin(x) : +\omega.$
- $penguin(TWEETY) : +\omega.$

As we see, the relative magnitude of priority constants in the heads of the rules captures the intuition that the default conclusion obtained by applying the second rule takes precedence over the conclusion obtained by the first rule. Note that, as the name suggests, priority constants represent the relative strength of evidences. Thus, the absolute magnitude of priority constants has no special significance, and only the relative aspect counts. Also, the relative strength of evidence that enables the body has no bearing on the relative strength of evidence in the conclusion, in the case of default rules. This is because, we are modelling defeasible reasoning, where at each step, we hold on to the most plausible conclusion, and where bodies of different rules are considered as independent sources of support, whose relative strengths are determined by the priority constants in rule heads. Note that these aspects distinguish our formalization from probabilistic approaches to inheritance [Pearl, 1988].

The priority constants are not restricted to be numerical values. One may in fact use constants representing the nodes of the network as priorities, with the information-ordering specified by the network topology. This is illustrated below by rewriting the first two defeasible rules of the Tweety example as follows, where $\perp$ represents the least positive evidence.

- $fly(x) : +bird \Leftarrow bird(x) : +\perp.$
- $fly(x) : -penguin \Leftarrow penguin(x) : +\perp.$

The ordering $<_k$ is set-up in such a way that the evidence-content of *penguin* is stronger than that of *bird,* which can be inferred from the network, because there is a directed path from *penguin* to *bird.*

Our theory allows representation of both strict and defeasible links. To specify that some of the rules are exception-free, we may add rules of the form: $p : -\omega \Leftarrow q : +\omega, p : +\omega \Leftarrow q : -\omega$ etc. These rules let guaranteed conclusions to propagate as such, which can be used to override conflicting default conclusions.

If the body of the rule is true by default, then the conclusion drawn from this rule must only have the status of a default conclusion. Priorities are used essentially only in arbitrating between conflicting conclusions; their absolute values are insignificant. The exception-free rules let us propagate and derive new conclusions that are known to be true or false beyond doubt. This is illustrated by the dead bird example:

- $dead(DODO) : +strict.$
- $bird(DODO) : +strict.$
- $fly(x) : -strict \Leftarrow dead(x) : +strict.$
- $fly(x) : +default \Leftarrow bird(x) : +strict.$

Even though there are conflicting conclusions about the flying ability of Dodo, the conflict can be resolved in favour of Dodo being unable to fly by virtue of being dead. This is because the latter conclusion is strict, as opposed to the default conclusion through bird.

To illustrate the treatment of ambiguity, we consider the Nixon diamond example.

- $republican(NIXON) : +\omega.$
- $quaker(NIXON) : +\omega.$
- $pacifist(x) : -1 \Leftarrow republican(x) : +1.$
- $pacifist(x) : +1 \Leftarrow quaker(x) : +1.$

The unique supported model associated with this specification contains *pacifist(NIXON)* : *1, meaning that our knowledge about Nixon being a pacifist is inconsistent.

Our formalization allows inheritances through paths containing negative arcs.

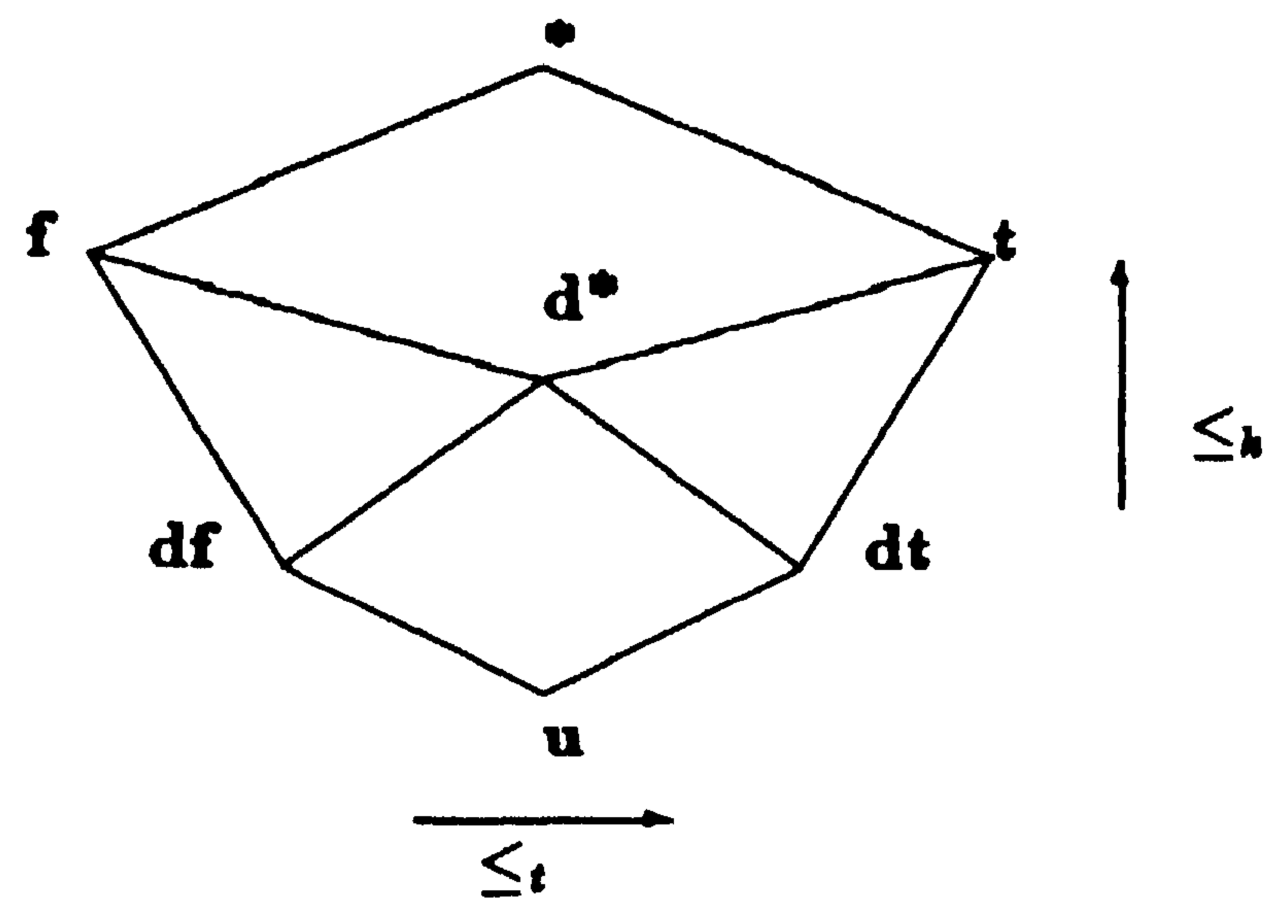

Figure 2: Ginsberg's bilattice for default reasoning

- $Pro\_American(GBR) : +\omega$.
- $Pro\_American(INDIA) : +\omega$.
- $Pro\_Soviet(CUBA) : +\omega$.
- $Pro\_Soviet(INDIA) : +\omega$.
- $Vote(x) : +1 \Leftarrow Pro\_American(x) : +\omega$.
- $Vote(x) : -1 \Leftarrow Pro\_Soviet(x) : +\omega$.
- $Support\_Contras\_Rebels(x) : +1 \Leftarrow Vote(x) : +1$.
- $Support\_Sandinistas(x) : +1 \Leftarrow Vote(x) : -1$.

One may interpret "+" as a "yes"-vote, "—" as a "no"-vote and "*" as an abstention. From the above specification we conclude that GBR supports the Contra rebels, while Cuba supports the Sandinistas. We cannot offer any conclusions about India's intentions because it abstains from voting.

There are advantages over the theory presented in [Horty *et* a/., 1987] that accrue out of ordering evidences to resolve ambiguities that are not inherent to the problem. This is illustrated by the following example.

- $toy(JAY) : +\omega$.
- $bird(JAY) : +\omega$.
- $fly(x) : -2 \Leftarrow toy(x) : +\omega$.
- $fly(x) : +1 \Leftarrow bird(x) : +\omega$.

Even though the topology of this network resembles the one for Nixon diamond, we may resolve the ambiguity in favour of toys being unable to fly, while [Horty *et al.,* 1987] cannot differentiate between these two cases.

The notion of supportedness buys us certain advantages over [Ginsberg, 1986] as illustrated below.

- *bird.* (Fact)
- *bird* $\Rightarrow$ *flies.* (Defeasible rule)
- *flies* $\rightarrow$ $\neg$*acro.* (Strict rule)

According to [Ginsberg, 1986J, this admits twc $\sqsubseteq_k$-minimal models, where the truth-lattice and information-lattice corresponding to these values is given in Figure 2. In one model *flies* is true by default and *aero* is false by default, while in the other *flies* is false, while *aero* is unknown.

Ginsberg notes that this does not allow one to conclude that the bird flies by default, because such a conclusion is not true in all $\sqsubseteq_k$-minimal models. He goes on to propose a criterion to prefer the first model over the second one by formalizing a notion of "strength of assumption" and discarding the second model by virtue of being based on "unreasonably" strong assumptions compared to the first one. We, on the other hand, do not run into such a problem. Indeed, in our formalism we would represent that problem as follows:

- $bird : +\omega$.
- $flies : +1 \Leftarrow bird : +1$.
- $acro : -\omega \Leftarrow flies : +\omega$.
- $acro : -1 \Leftarrow flies : +1$.

This admits a unique supported model containing *flies: +1* and *acro: —1* . The interpretation corresponding to the second minimal model in the Ginsberg's approach is not supported because *flies: —w* (i.e., *flies* is false) does not have any justification.

One may try to translate inheritance networks into our logic to give them a formal semantics. The main problem here is the choice of priority constants based on the network topology. We will sketch such an approach below.

Define the *evidence* set corresponding to a network to be $\{+,-\} \times (\{\perp,\omega\} \cup \mathbf{N})$. The first component of an evidence represents its truth-content, while the second component captures the information-content. repre-sents minimum information, while $\perp$ represents maximum information. The ordering of entities in N on the information scale is consistent with the -<-ordering of the corresponding nodes in the network. That is, if p is a subclass of q $(p \prec q)$, then $+q <_k +p$ and — *q* $<_k$ -fp (similarly for — *p).*

We choose *priority constants* to represent finite subsets of the evidence set. Informally, a priority constant specifies the set of evidences in support of an atom. The truth and the information-content of a priority constant, 7r, is determined solely by the $<_k$-maximal elements of IT. For instance, in the Tweety example, the priority constant corresponding to *fly(TWEETY)* is *{-fbird, —penguin},* which is equivalent to *{—penguin}.* The *lubu* of a set of priority constants is obtained by taking their union. For instance, in the Nixon diamond example, the *lubk* of *{-republican}* and *{-i-quaker}* is *{-republican,+quaker}.* Because *-republican* and *-hquaker* are incomparable on the information scale and they have conflicting truth values, the conclusion about Nixon's pacifism is inconsistent.

The translation of a positive arc from an individual node i to a property node p is: $p(i) : \{+\omega\}$; a negative arc is translated as: $p(i) : \{-\omega\}$. The translation of a positive arc from a property node p to a property node q is: $p(x) : \{+q\} \Leftarrow q(x) : \{+\perp\}$; the translation of a negative arc is: $p(x) : \{-q\} \Leftarrow q(x) : \{+\perp\}$. The meaning assigned to an inheritance network corresponds to the unique supported model associated with its translation.

An advantage of our approach is that inheritance specifications can be designed in the same style as Prolog programs. This is because the meaning of a predicate depends only on the meanings assigned to the predicates in bodies of rules defining that predicate.

4 Conclusion

We have described a logical framework for a theory of inheritance obtained by a novel combination of concepts from logic programming and multi-valued logics. A Horn-clause logic language is used to specify the networks. A model theory for the language is provided and a unique supported model is associated with each network. We have also presented an algorithm to compute this model. Our theory resembles the skeptical theory of [Horty *et* a/., 1987], but is more general. "Preferential" inheritance and inheritance through paths containing negative arcs can be expressed in our formalism. Our framework can also be extended to formalize credulous theories of inheritance.

A cknowledgements

We would like to thank David S. Warren for enlightening discussions on the subject of inheritance and the referees for their comments and suggested improvements.

References

[Apt *et* al, 1987] K. R. Apt, H. A. Blair and A. Walker, Towards a theory of declarative knowledge, *Foundations of Deductive Databases and Logic Programming,* Jack Minker(ed.), Morgan Kaufman, pp 89-148, 1987.

[Belnap, 1977a] N. Belnap, How a computer should think, *Contemporary Aspects of Philosophy,* G. Ryle(ed.), Oriel Press, pp 30-56, 1977.

[Belnap, 1977b] N. Belnap, A useful four-valued logic, *Modern Uses of Multi-valued Logic,* J. Dunn and G. Epstein (ed.), D. Reidel, pp 8-37, 1977.

[Blair, 1987] H. A. Blair and V. S. Subramanian, Paraconsistent Logic Programming, *Lecture Notes in Computer Science,* vol. 287, pp 340-360, 1987.

[Etherington, 1983] D. Etherington and R. Reiter, On inheritance hierarchies with exceptions, In *Proceedings of the Second National Conference on Artificial Intelligence,* pp 104-108, 1983.

[Fitting, 1988] M. Fitting, Bilattices and the Semantics of Logic Programming, To appear in *Journal of Logic Programming,*

[Ginsberg, 1986] M. Ginsberg, Multi-valued logics, Technical Report 86-29, Department of Computer Science, Stanford University, 1986.

[Haugh, 1988] B. Haugh, Tractable Theories of Multiple Defeasible Inheritance in Ordinary Nonmonotonic Logics, In *Proceedings of the Seventh National Conference on Artificial Intelligence,* pp 421-426, 1988.

[Horty *et al,* 1987] J. Horty, R. Thomason, and D. Touretzky, A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks, In *Proceedings of the Sixth National Conference on Artificial Intelligence,* pp 358-363, 1987.

[Horty and Thomason, 1988] J. Horty, R. Thomason, and D. Touretzky, Mixing Strict and Defeasible Inheritance, In *Proceedings of the Seventh National Conference on Artificial Intelligence,* pp 427-432, 1988.

[Kifer and Li, 1988] M. Kifer, and A. Li, On the Semantics of Rule Based Expert Systems with Uncertainties, In *Proceedings of the Second International Conference on Database Theory,* Bruges, Belgium, 1988.

[Kifer and Lozinskii, 1989] M. Kifer, and E. L. Lozinskii, RI: A Logic for Reasoning with Inconsistency, In *Proceedings of Fourth IEEE Symposium on Logics in Computer Science,* June 1989.

[Krishnaprasad *et* a/., 1988a] T. Krishnaprasad, M. Kifer, and D. S. Warren, A Circumscriptive Semantics of Inheritance Networks. (Submitted)

[Krishnaprasad and Kifer, 1988b] T. Krishnaprasad, M. Kifer, and D. S. Warren, On the Declarative Semantics of Inheritance Networks, In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence,* August 1989.

[Moore, 1985] R. M oore, Semantical Considerations on Nonmonotonic Logic, *Artificial Intelligence 13*, pp 75-94, 1985.

[Padgham, 1988] L. Padgham, A Model and Representation for Type Information and Its Use in Reasoning with Defaults, In *Proceedings of the Seventh National Conference on Artificial Intelligence,* pp 409-414, 1988.

[Pearl, 1988] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufman, San Mateo, 1988.

[Przymusinska and Gelfond, 1988] H. Przymusinska and M. Gelfond, Inheritance hierarchies and autoepistemic logic, Technical Report, Department of Computer Science, University of Texas at El Paso, 1988.

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

[Touretzky, 1986] D. Touretzky, The mathematics of Inheritance Systems, Morgan Kaufmann, Los Altos, 1986.

[Touretzky, 1987] D. Touretzky, J. Horty, and R. Thomason, A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems, *In Proceedings of the Tenth International Joint Conference on Artificial Intelligence,* Milan, pp 476-482, 1987.