On the Declarative Semantics of Inheritance Networks

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

Abstract

Usually, semantics of inheritance networks is specified *indirectly* through a translation into one of the standard logical formalisms. Since such translation involves an *algorithmic* aspect, which is usually complex, these approaches to inheritance are not truly declarative. We provide a general framework for specifying a *direct* semantics of inheritance networks. Because the networks are not expressive enough to capture all intuitions behind inheritance, a number of significantly different semantics have been proposed. Our approach allows us to give direct semantics to a number of different proposals found in the literature, and clarifies the relationships among them. It also provides a yardstick for measuring adequacy of translation into logical formalisms of various intuitions about inheritance.

1 Introduction

Inheritance networks represent individuals, classes and properties. For efficient representation and determination of properties of individuals, these networks have evolved from simple property lists to class-subclass hierarchies to multiple inheritance networks.

Any realistic representation of the real-world knowledge must necessarily allow representation of exceptions. For instance, typically mammals are nonflyers. Bats are exceptional mammals that normally fly. But dead bats do not fly. In general, inheritances from subclasses must dominate over inheritances from classes in the case of a conflict.

The representation language should also allow expression *of preferential* inheritance of a property from a class over inheritance from another class. For instance, given that man is an omnivore, that is, man is both a herbivore and a carnivore, one must be able to infer that men have canines. Even though herbivores typically do not possess canines, man "inherits" them because he is a carnivore.

Providing a satisfactory semantics to nonmonotonic multiple inheritance networks poses a significant intellectual challenge, and a number of different proposals have appeared in the literature. Normally, the semantics is given through a translation into a logical formalism as in [Etherington, 1983], [Ginsberg, 1988], [Haugh, 1988], [Krishnaprasad *et a*/., 1988a], [Krishnaprasad and Kifer, 1988b], [Przymusinska and Gelfond, 1988], or by developing special purpose techniques as in [Horty *et a*/., 1987], [Padgham, 1988], [Touretzky, 1986]. However, none of these approaches can be regarded as truly declarative. The problem is that an *algorithmic* transformation from inheritance networks to some logical formalism is given, and only then its semantics described in terms of this formalism. Depending on the conceptual difficulty of the algorithmic part, the claim of "declarativeness" of each specific approach is unfounded to a different degree. In addition, these approaches are usually complex and it is not clear how the different semantics relate to each other.

In this paper, we provide a general framework for specifying declarative semantics of inheritance networks *directly.* From the users point of view, semantics of a network should be given directly in terms of the network and be easily comprehensible the same way the semantics of logic programming is. Because the networks are not expressive enough to capture differing intuitions behind inheritance, a number of significantly different semantics have been given. Our unified approach permits a careful study of the differences among some of the extant proposals for a theory of inheritance and it provides a yardstick for measuring the correctness of a translation and implementation of various semantics.

In Section 2, following [Krishnaprasad *et al.,* 1988a], we present a syntax of inheritance networks that generalizes the earlier definition of [Touretzky, 1986] allowing representation of the above features.

Section 3 motivates through examples different semantics of inheritance networks stemming from differing interpretations of ambiguity. (See Figure 2.) In particular, we specify constraints that a semantic structure must satisfy to be a model of the network. We notice that different approaches differ in the specification of these constraints. We also present minimality conditions for selecting certain "preferred models", which can be assumed as representing the meaning of the network.

Section 4 discusses the fundamental semantic differences between the bottom-up view of inheritance [Haugh, 1988] [Horty *et a*/., 1987] [Krishnaprasad *et a*/., 1988a] [Krishnaprasad and Kifer, 1988b] and the top-

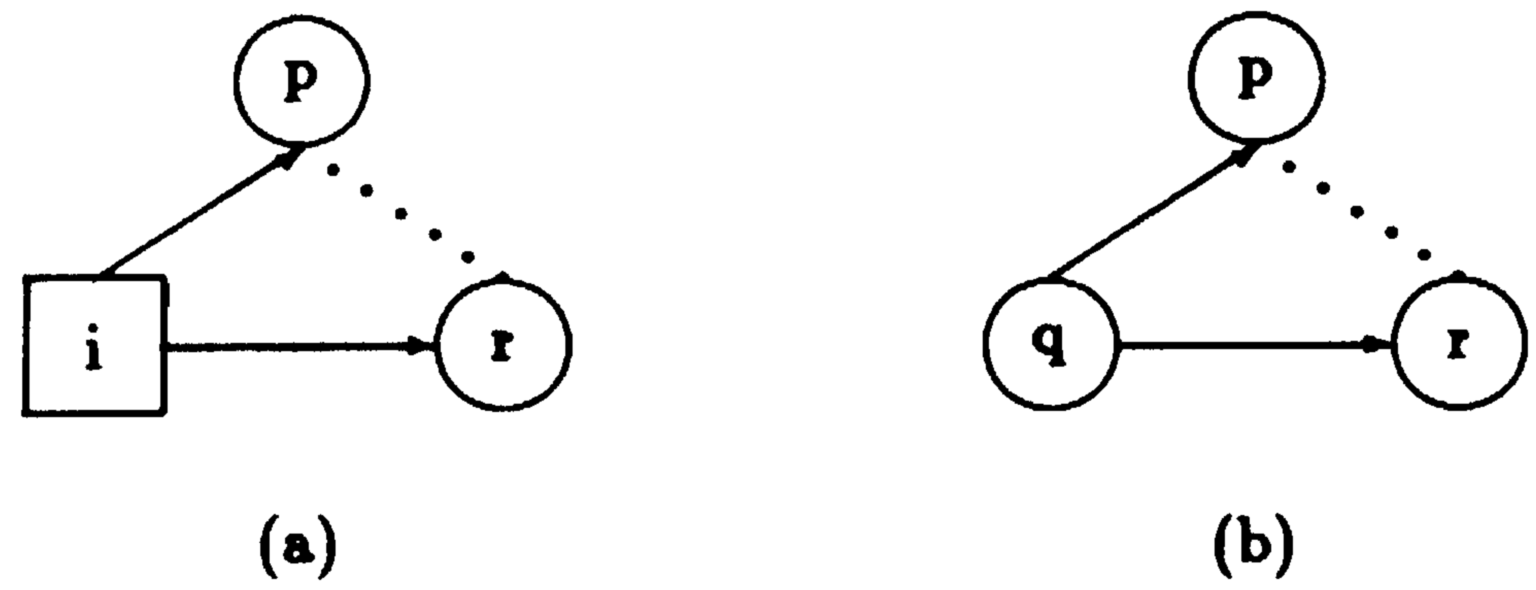

Figure 1: Constraints on pre-order $\preceq_p$

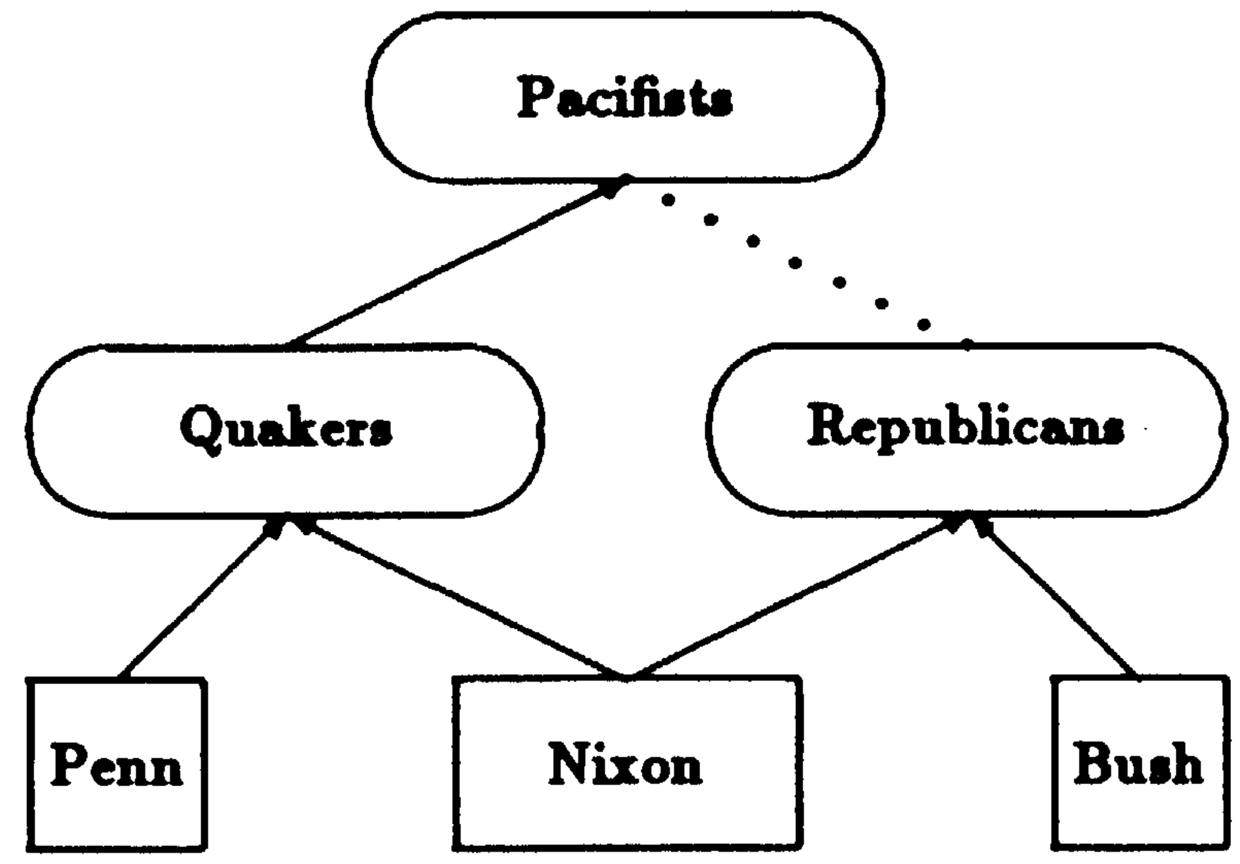

Figure 2: Nixon diamond network

down view espoused in [Touretsky, 1986].

2 Syntax

A *preferential inheritance network* is an ordered directed acyclic graph consisting of a set of individual nodes **I**, a set of class/property nodes **N**, a set of positive arcs $\mathbf{E}^+ \subseteq (\mathbf{I} \cup \mathbf{N}) \times \mathbf{N}$, a set of negative arcs $\mathbf{E}^- \subseteq (\mathbf{I} \cup \mathbf{N}) \times \mathbf{N}$, and for each node $p \in \mathbf{N}$, a pre-order relation $\preceq_p$ on its in-arcs. We use p, q and r to stand for any property node in **N**, and i for any individual node in **I**. Assume $\mathbf{E}^+ \cap \mathbf{E}^- = \emptyset$. Let $e \sim f$ stand for ($e \preceq f \wedge e \preceq f$), and $e \prec f$ for ($e \preceq f \wedge e \not\preceq f$).

An arc $\langle i, p \rangle$ stands for the assertion – "i is a p" and an arc $\langle p, q \rangle$ stands for the defeasible rule – "Normally, p's are q's".

The pre-order relation $\preceq_p$ [Krishnaprasad *et al.*, 1988a] captures the relative strength of inheritance through the in-arcs of p, so it must satisfy the following conditions:

1. An explicitly stated assertion "i is a p" must have priority over the default inheritance of $\neg p$ by i through r. (See Figure 1(a).) That is, for arcs $\langle i, p \rangle$ and $\langle r, p \rangle$, we have $\langle r, p \rangle \prec \langle i, p \rangle$.
2. For all individuals in q, inheritance of p ($\neg p$) through a subclass q must take precedence over inheritance of $\neg p$ (p) through the class r. (See Figure 1(b).) That is, for any pair of arcs $\langle q, p \rangle$ and $\langle r, p \rangle$, such that there is a directed path from q to r, it is the case that $\langle r, p \rangle \prec \langle q, p \rangle$.

3 Semantics

To specify the states of the world a network represents, we associate a set of models with it. We define the notion of a semantic structure for a network and provide constraints for the structure to be a model. These constraints embody intuitions about how individuals "move up" the network inheriting properties on the way.

The differences between the various semantics of inheritance networks found in the literature stem from differing intuitive understandings of what a network with a topology similar to that in Figure 2 means. We motivate and specify different direct semantics to inheritance networks.

3.1 Skeptical Semantics

3.1.1 Motivation

Consider Figure 2 representing the following facts. Quakers are normally Pacifists, while Republicans are not. Nixon is both a Quaker and a Republican. Penn is a Quaker and Bush is a Republican. The skeptical semantics supports the conclusion that Penn is a pacifist, while Bush is not. But it is ambiguous about Nixon's pacifism. This is legitimate because we have two equally strong conflicting evidences in support of his pacifism and there is no way to choose one over the other to determine whether Nixon is a pacifist or not.

3.1.2 Skeptic Models

A *semantic structure* for a network is an assignment of a distinct individual constant n for each node n in I and a triple of sets of individuals (p+,p*,p-) for each node p in N. Intuitively, p^+ contains individuals known to have property p, *p~* contains individuals known to have property -p, and p* contains individuals which are inconclusive about p, i.e., those which are both in p^+ and p-.

We use the letters n, m to denote individual nodes in I and n, m to denote the respective individuals assigned to these nodes. The letters p, q, r, s, t will be used to denote property nodes in N. We also assume, for notational convenience, the equivalences: $p^+(n) \equiv n \in p^+$ (resp. p^-, p^*) and $\neg p^+(n) \equiv n \notin p^+$, (resp. p-,p*).

The constraints for a semantic structure to be a model are given below. Informally, n should inherit p if the maximal evidence in support of the inheritance is through a positive arc to p; n inherits -p if the maximal evidence is through a negative arc to p; and n is ambiguous about inheriting p or -p, whenever there are equally strong or incomparable evidences for p and -p.

The statement that maximal evidence for -p(n)" can be formally expressed as:

$$\mathbf{E}^-(n,p) \ \vee \ \exists t \, [\, \mathbf{E}^-(t,p) \wedge t^+(n) \wedge$$
$$\neg \exists s \, (\, [\mathbf{E}^+(s,p) \wedge s^+(n)\,] \wedge \langle t,p \rangle \prec \langle s,p \rangle \,) \,],$$

and the statement that "∃ *maximal evidence for p(n)*" can be formally expressed as:

$$\mathbf{E}^+(n,p) \;\vee\; \exists t\,[\,\mathbf{E}^+(t,p) \wedge t^+(n) \wedge$$
$$\neg\exists s\,(\,[\,\mathbf{E}^-(s,p) \wedge s^+(n)\,] \wedge \langle t,p\rangle \prec \langle s,p\rangle\,)\,].$$

For each property node p and individual node n :

$$p^-(n) \quad \text{if} \quad \text{"}\exists \textit{ maximal evidence for } \neg p(n)\text{"}. \quad (1)$$
$$p^+(n) \quad \text{if} \quad \text{"}\exists \textit{ maximal evidence for } p(n)\text{"}. \quad (2)$$
$$p^*(n) \;\equiv\; p^+(n) \wedge p^-(n).$$

3.1.3 Minimal Models

We may select "preferred" models, by capturing the idea of minimality, by replacing if with iff in (1) and (2). The meaning so obtained corresponds to the skeptical theory of inheritance given in [Krishnaprasad and Kifer, 1988b].

Note that checking whether a semantic structure is a model or a minimal model is local, i.e., for each node it depends only on its immediate descendents.

3.2 An alternative skeptical semantics

The skeptical theory of [Horty *et a*/., 1987] is obtained from the above semantics by associating with each node a pair of sets (p^+_{hor}, p^-_{hor}), where

$$p^+_{hor} = p^+ - p^* \text{ and } p^-_{hor} = p^- - p^*.$$

The semantics of [Horty *et a*/., 1987] does not have p*, and treats "no information" and "ambiguous information" on par. In contrast, [Krishnaprasad and Kifer, 1988b] distinguishes between these two situations.

3.3 Credulous Semantics

3.3.1 Motivation

Consider again Figure 2. The skeptical semantics given above places Nixon in *pacifist**, declaring that the information about Nixon being a pacifist is ambiguous. On the other hand, it is equally reasonable to argue that in real life Nixon is either a pacifist or a non-pacifist, and to associate a set of two minimal models with this network to capture this situation [Touretzky, 1986]. In one model, Nixon is a Pacifist, while in the other, he is not.

3.3.2 Credulous Models

A *semantic structure* for a credulous theory is different from that of a skeptical theory only in that there is no set p* associated with property nodes. In other words, a semantic structure assigns a pair of sets of individuals (P^+, P^-) *to each* node p in N.

The constraints that a semantic structure must satisfy to be a model are given below. Informally, n inherits p if the maximal evidence in support of the inheritance is through a positive arc to p, n inherits -p if the maximal evidence is through a negative arc to p. In the case that there are equally strong or incomparable evidences to inherit both p and -p, n inherits p in one model and inherits -p in the other.

For each property node p and individual node n :

$$p^-(n) \text{ if } \neg p^+(n) \wedge \text{"}\exists \textit{ maximal evidence for } \neg p(n)\text{"}. \quad (3)$$
$$p^+(n) \text{ if } \neg p^-(n) \wedge \text{"}\exists \textit{ maximal evidence for } p(n)\text{"}. \quad (4)$$
$$p^+ \cap p^- \;=\; \emptyset. \quad (5)$$

Observe that credulous semantics differs from the skeptic one essentially in the first conjunct of (3) and (4) (cf. (1) and (2)). Note that checking whether a semantic structure is a model is again local.

3.3.3 Minimal models

We may filter out "extraneous" models by capturing minimality, similarly to the skeptic case by replacing if with iff in (3) and (4). In this case, (5) would follow from the modified (3) and (4).

This semantics gives two minimal models for the Nixon diamond network in Figure 2. In one model, Nixon is an abnormal Republican inheriting Pacifism from Quakers, while in the other model, he is an abnormal Quaker and a non-Pacifist. This specifies the credulous semantics capturing off-path preemption of [Haugh, 1988] and the credulous evidence-based semantics obtainable along the lines of [Krishnaprasad and Kifer, 1988b].

3.4 An alternative credulous semantics

Consider once again Figure 2. In addition to the conclusions sanctioned by the above approach, we may conclude further that Bush is not a Quaker and Penn is not a Republican. The circumscriptive semantics presented in [Krishnaprasad *et a*/., 1988a] captures this intuition to strengthen negative conclusions. We specify this by modifying the constraint (3) of Section 3.3.2 and augmenting it with two disjuncts as follows:

$$p^-(n) \quad \text{if}$$
$$\neg p^+(n) \;\wedge\; [\;\text{"}\exists \textit{ maximal evidence for } \neg p(n)\text{"}$$
$$\vee\; \exists t\,(\,\mathbf{E}^+(p,t) \wedge t^-(n) \;\vee\; \mathbf{E}^-(p,t) \wedge t^+(n)\,)\,]. \quad (6)$$

Furthermore the minimality constraint is obtained by changing if to iff.

Note that checking whether a semantic structure is a model or a minimal model is still local.

3.5 Ambiguity-propagating Skepticism

One may associate a unique meaning with each inheritance network by taking the intersection of the minimal models sanctioned by the semantics in Section 3.3.3. This is the ambiguity-propagating version of skepticism of [Touretzky *et a*/., 1987].

3.6 Generic Statements

The meaning of a node can be defined either as a set of individuals or as a set of properties. To determine the set of *properties* that may be associated with a "typical" individual of a class by virtue of being a member of that class, imagine associating with each class node a special individual node as its child. This is achieved by extending the domain of discourse from the set of constants representing individuals in I to the set of constants representing all nodes in I U N, and adding the constraint

$P \in P^+$, for each node p in N. The properties inherited by the constant p corresponding to node p gives the collection of properties associated with a "typical" member of p. These inherited properties can also be interpreted as the set of generic statements supported by our network.

3.7 Mono tonic Inheritance

It is also easy to specify the approach of [Thomason *et* a/., 1987] to monotonic semantic networks in the present framework by interpreting each arc as representing a generic statement that does not admit any exception. In the case of a conflict, an inherited property may be assigned the status of inconsistent belief. We can specify the constraints on an interpretation as follows:

$$p^+(n) \text{ iff } E^+(n,p) \vee \exists t\, [\, E^+(t,p) \wedge t^+(n)\,].$$

$$p^-(n) \text{ iff } E^-(n,p) \vee \exists t\, [\, (\, E^-(t,p) \wedge t^+(n)\,) \vee (\, E^-(p,t) \wedge t^+(n)\,) \vee (\, E^+(p,t) \wedge E^-(n,t)\,)\,].$$

The meaning assigned to each node is obtained by associating with each node a triple of sets $(p^t, p^f, p^\top)$, where

$$p^\top = p^+ \wedge p^-,$$

$$p^t = p^+ - p^\top \text{ and } p^f = p^- - p^\top.$$

4 Discussion

Note that ascertaining whether a preferential network is syntactically well-formed (see conditions 1 and 2 of Section 2) may involve nonlocal inspection of the entire subnet connecting the descendents of a node. However, computing the meaning of a node is *local* in that it depends only on the meanings assigned to its immediate neighbors and not on the meanings assigned to nodes arbitrarily far away. The importance of locality in inheritance networks was also observed in [Ginsberg, 1988]. On the other hand, the semantics described in [Touretzky, 1986] is nonlocal which is partly responsible for its computational complexity.

The locality aspect highlights the differences between the two prevailing views of inheritance: One view of inheritance is that the individuals move up the network to acquire properties. Another view of inheritance is that the properties flow down the network to apply to the individuals [Touretzky *et* a/., 1987]. We will refer to the former as the *primal* view, and to the latter as the *dual* view. Similar to the semantics for the primal view presented in this paper, it is possible to associate a set of properties as the dual meaning of a node. Both these views are "equivalent" and easy to characterize for exception-free inheritance networks or for tree-structured class-property hierarchies. However, this duality breaks down for inheritance networks with multiple inheritance and exceptions. As a matter of fact the primal and dual semantics have fundamental differences and it turns out that it is easier to capture one's intuitions about inheritance using the primal meaning of a network. This is illustrated in [Krishnaprasad *et* a/., 1988a] by comparing the circumscriptive theory of inheritance given, to the formalization in [Touretzky, 1986]

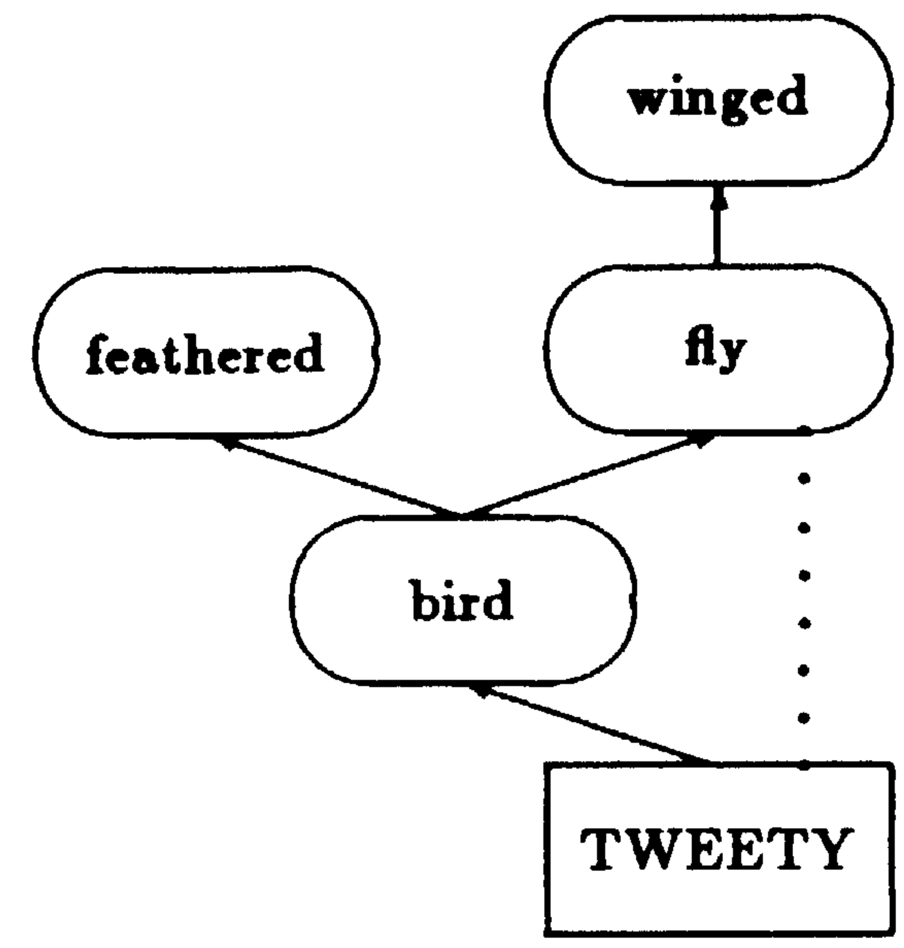

Figure 3: Wings Example

which tries to characterize the dual meaning. In particular, the latter leads to a loss of locality as illustrated by the example of Figure 3.

All the semantics presented above associate a unique minimal model with the network in Figure 3. In this model, Tweety is a flightless feathered bird, but it is not known whether Tweety is winged.

The dual meaning of the network assigns the set { *fly, winged, feathered* } to the node *bird* and the set { *-fly, bird, feathered* } to the node *Tweety.* One can conclude that Tweety is feathered because it is a bird, but not that it is winged. Birds are winged because they fly. But Tweety does not fly, hence it should not inherit wings just by virtue of being a bird. Thus, the meaning of the node *Tweety* cannot be determined by looking only at the meanings of its immediate ancestors. Instead, we need to know the context (in the form of inheritance paths) to see whether the property of wingedness is inherited or not. Thus, the reason for the difficulties with the dual semantics is that in order to obtain an adequate formalization of inheritance, it seems necessary to sacrifice the afore-mentioned locality of the semantics, as it is done in [Touretzky, 1986]. Such formalization necessarily requires considering paths in the network, which means that the dual direct semantics similar to the one presented here will be rather awkward. For a similar reason, the coupling of individuals, as in [Touretzky, 1986], brings nonlocality in.

In characterizing the dual meaning of a network, we cannot incorporate preferential inheritance smoothly. (This will be clear from the following example.) This is a limitation because there are a number of natural cases where the topology resembles the one for an ambiguous network, but there is no inherent ambiguity in the problem. Many a times the ambiguity may be due to incompleteness in our knowledge, and this may get resolved as the network evolves. For instance, consider an example from the medical expert system [Borgida, 1988]. A patient suffering from renal failure has high blood pressure, while a hemorrhaging patient has low blood pressure. If we have a hemorrhaging patient with renal failure, the above facts do not allow us to conclude whether the patient has high blood pressure or low blood pressure. But then it is well-known that the patient would have low blood pressure due to loss of blood. Preferential inheritance allows expression of such meta-knowledge. In particular, one may specify how the ambiguity is to be resolved for every pair of nodes with respect to a property. There does not seem to be a natural way of determining dual preferential semantics without compromising locality.

5 Limitations and Possible Extensions

The examples in Section 3 deal only with *homogeneous* inheritance networks [Touretzky *et* al., 1987], that is, there is no explicit distinction between default and strict arcs. It is possible to extend the syntax and specify appropriate constraints to capture the semantics presented in [Krishnaprasad and Kifer, 1988b] for *heterogenous* networks [Touretzky *et* a/., 1987] and the one in [Etherington, 1983] for networks with *explicit* exception links. The semantics of heterogenous networks of [Horty *et* a/., 1987] and the semantics of homogeneous networks of [Touretzky, 1986] are intrinsically *nonlocal* and are unlikely to have elegant representation in our framework.

6 Conclusion

We thus have a formal general framework to describe direct semantics to inheritance networks and to derive set-based bottom-up inheritance algorithms from it. We can also compare different proposals found in the literature by specifying them in this framework. This also emphasizes the importance of locality in the specification of the semantics and sheds light on the differences in the primal and the dual semantics of the network.

Acknowledgements

We would like to thank Weidong Chen for fruitful discussions on the subject of this paper and the referees for their helpful comments.

References

[Borgida, 1988] A. Borgida, Modeling Class Hierarchies with Contradictions, In *Proceedings of SIGMOD'SB Conference,* pp 434-443, 1988.

[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.

[Ginsberg, 1988] M. L. Ginsberg, A Local Formalization of Inheritance: Preliminary Report, 1988. (Manuscript)

[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* a/., 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.

[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 and M. Kifer, An Evidence-based Framework for a Theory of Inheritance, In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence,* August 1989.

[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.

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

[Thomason *et* a/., 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 *et* a/., 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,* pp 476-482, 1987.