Information Processing Letters 46 (1993) 263–268
Elsevier
26 July 1993

Locality in inheritance networks *

Krishnaprasad Thirunarayan
Department of Computer Science and Engineering, Wright State University, Dayton, OH 45435, USA

Communicated by D.A. Plaisted
Received 11 March 1992
Revised 26 April 1993

Keywords: Representation and reasoning; semantics and algorithms

1. Introduction

Touretzky proposed a rigorous theory of nonmonotonic inheritance in [11] which was later found to be intractable [9] and to lead to counter-intuitive interpretation of certain networks [12]. Subsequently, there have been a number of proposals for semantics of inheritance networks – done either through a translation into a general logical formalism [1,6,8], or by special path-based techniques [2,3,11].

As an alternative to these, we designed *local* theories of inheritance that are conceptually simple and computationally attractive [5]. Here, we formalize *locality* as it applies to inheritance networks and illustrate it with examples. We then generalize the notion of local semantics to what we call ground local semantics to further understand and clarify the relationship between our theories and the path-based theories. Our theories are both local and ground local; the skeptical semantics of [2] is not local, but is ground local, while the semantics of [11] is neither local nor ground local.

2. A framework for semantics specification

A *preferential network* is an ordered DAG consisting of individual nodes **I**, property nodes **P**, positive arcs [1] $\mathbf{E}^+ \subseteq (\mathbf{I} \cup \mathbf{P}) \times \mathbf{P}$, negative arcs $\mathbf{E}^- \subseteq (\mathbf{I} \cup \mathbf{P}) \times \mathbf{P}$, and for each node $p \in \mathbf{P}$, a local specificity relation $\prec_p$ on its *in-arcs*. Let p, q, r stand for a node in **P**, and i, m, n for a node in **I**.

A (Herbrand) *semantic structure* for a preferential network is a pair $(\mathscr{D}, \mathscr{M})$, where $\mathscr{D}$ is the set of nodes and $\mathscr{M}$ is the set-based meaning function [6]. $\mathscr{M}$ associates with each node $p \in \mathbf{P}$ a pair of subsets (p^+, p^-) of $\mathscr{D}$, where p^+ (resp. p^-) represents the individuals/properties that "inherit" property p (resp. $\neg p$).

The different approaches to inheritance differ in the specification of the constraints for a semantic structure to be a *model*. Informally, individual n should inherit property p (resp. $\neg p$) if the maximal evidence in support of the inheri-

Correspondence to: T. Krishnaprasad, Department of Computer Science and Engineering, Wright State University, Dayton, OH 45435, USA. Email: tkprasad@cs.wright.edu.

* This research was supported in part by the NSF grant IRI-9009587.

[1] An arc from an individual node to a property node represents an *instance-of* relation, while an arc from a property node to a property node represents an *is-a* relation.

tance is through a positive (resp. negative) arc to p. All the inheritance theories developed in [5,6] can be specified using local constraints of the form: For each property p and individual/property n,

$$p^-(n) \textbf{ if } BODY \quad p^+(n) \textbf{ if } BODY,$$

where the BODY [2] formula refers only to the meaning assigned to p's immediate neighbors and not to the meaning assigned to nodes arbitrarily far away. Furthermore, we may filter out "extraneous" models by capturing minimality, by replacing **if** with **iff**.

The *local specificity* used to resolve inheritance conflicts is as follows: If, on the basis of the fact that x is a q, we can infer that x is a p, then we regard q as providing more specific information than p. The intuition is that if we associate a property r with q, then we are willing to override any evidence for $\neg r$ that rests on the "implicit" property p of the members of q. Determination of the local specificity $\prec_r$ can be integrated with the inheritance algorithm.

3. Locality in inheritance networks

Locality is fundamental to efficient representation and use of information. The notion of local semantics of inheritance networks can be formalized by making precise the statement – *the meaning of a node is constrained* only *by the meaning assigned to the nodes in its* immediate *neighborhood*. As a first cut, we can assert that *two nodes with "similar" neighborhoods should have the same meaning*. We consider two nodes p and q to have *similar* neighborhoods if: (i) the number of nodes adjacent to p and q are equal, and (ii) there is a one-to-one correspondence between the nodes adjacent to p and those adjacent to q such that (a) the corresponding nodes have the same meaning, and (b) the relative strength of inheritance through the corresponding arcs to p and q, are the same. (Consider the penguin triangle and the Nixon diamond networks. The nodes for *pacifism* in the Nixon diamond and *fly* in the penguin triangle do not have similar neighborhoods because the inheritance through the arc from *penguin* to *fly* is stronger than that from *bird* to *fly*, while the strength of inheritance through the arcs from *republican* and *quaker* to *pacifism* are incomparable. As a result, the penguin triangle is unambiguous, while the Nixon diamond is ambiguous.) However, requiring that p and q have identical meaning whenever they have similar neighborhoods, seems overly restrictive. This is because, an intuitively satisfactory assignment to a node can have "unsupported" redundant information that does not violate the imposed local constraints. (Given that birds fly and that Tweety is a bird, one can imagine two possible states of the world – in one, only Tweety can fly; while in the other, both Tweety and the Angel can fly. Here, we have no evidence for (or against) the Angel to fly.) To accomodate this interpretation, we weaken the requirement for locality as follows: *Two nodes with* similar *neighborhoods should have* interchangeable, *but not necessarily identical, meaning*. We make rigorous these ideas in the following definitions.

Let Γ be a preferential network. A node r is a *neighbor* of a node p if there is a direct arc between r and p. A *node-environment* Γ_p corresponding to the node p of Γ is the subnetwork of Γ consisting of the node p and all its neighbors. A *local specificity relation* with respect to a node p (denoted $\prec_p$) is a binary asymmetric relation among the arcs of the node-environment of p. We say that there exists an *isomorphism* between the node-environments Γ_p^1 and Γ_q^2 if there exists a bijection between the nodes of Γ_p^1 and Γ_q^2 such that the *adjacency relation* among the nodes and the *local specificity relation* among the arcs of the node-environments are *preserved*. For our purposes, a semantics of inheritance networks is an assignment of sets of nodes to the nodes of the networks that satisfy a set of local constraints $\mathscr{C}$. A satisfying assignment to the network Γ is denoted as $\mathscr{M}_\Gamma$, and the meaning that $\mathscr{M}_\Gamma$ assigns to a node p is denoted as $\mathscr{M}_\Gamma(p)$. The notation $\mathscr{M}_\Gamma[\mathscr{M}_\Gamma(p) \leftarrow \text{S}]$ stands for the assignment obtained from $\mathscr{M}_\Gamma$ by replacing $\mathscr{M}_\Gamma(p)$ with S.

[2] To understand the theme of this paper, the precise details of the semantics are not important.

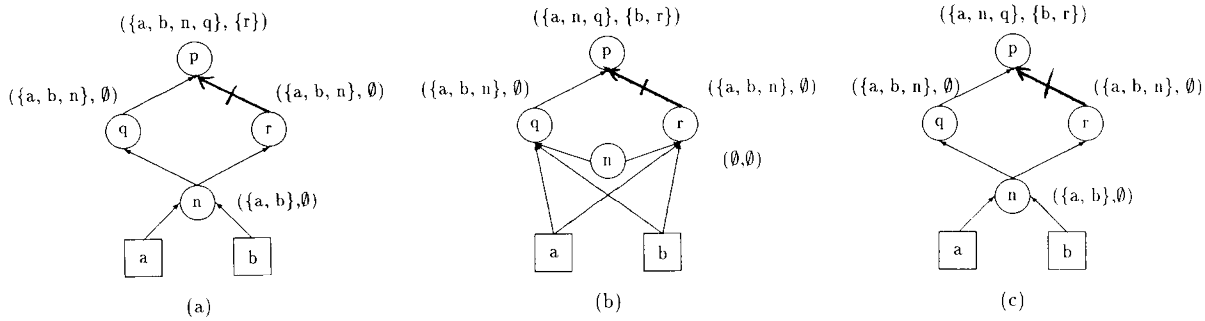

Fig. 1. Nonlocality of Touretzky's semantics I.

Definition 1. A semantics of inheritance networks is *local* if, for any pair of networks Γ^1 and Γ^2 containing isomorphic node-environments Γ_p^1 and Γ_q^2, whenever the meaning assigned to all the neighbors of node p in Γ_p^1 is identical to the meaning assigned to the "corresponding" neighbors of node q in Γ_q^2, it is the case that the assignment $\mathscr{M}_\Gamma^2[\mathscr{M}_\Gamma^2(q) \leftarrow \mathscr{M}_\Gamma^1(p)]$ satisfies Γ^2.

Theorem 2. *The semantics given in Section* 2 (*that is, in* [6]) *are* local.

We now show that the semantics developed in [11] is *not local* [3]. Figures 1(a) and 1(b) depict two networks and their meaning according to [11]. Each node p is assigned a pair of sets of nodes (p^+, p^-), where $q \in p^+$ stands for q possesses p and $q \in p^-$ for q possesses $\neg p$. Figure 1(c) depicts the application of the locality condition to p. If the semantics in [11] were local, the assignment given in Fig. 1(c) would satisfy the network. However, the assignment in Fig. 1(c) is not a model according to [11] because a and b do not inherit p similarly (that is, there is no "coupling" between a and b) [12]. Thus, the semantics can be nonlocal. In general, the *coupling* required in the downward inheritance reasoners makes this view of the semantics nonlocal.

But, as discussed in [4,12], one can claim that the chosen semantic structure is appropriate to capture only individual flow view of inheritance in upward reasoners, and that the "dual" semantic structure is needed to capture adequately property flow view of inheritance in downward reasoners. This change of view can, in fact, make the semantics of a network local. For instance, the semantic structure can be redefined to reflect the dual view [4] as follows: let the assignment of sets (p^+, p^-) to node p mean the following: $q \in p^+$ stands for p possesses q and $q \in p^-$ for p possesses $\neg q$. In that case, one can verify that the satisfying assignments to the networks in Fig. 1 do not violate locality. But consider Fig. 2. Figure 2(a) displays a model of the network in which it is not known whether t is a w. The assignment in Fig. 2(b) is a model of the modified network in which t is a w by virtue of being a member of b. However, the assignment in Fig. 2(c), obtained by the application of the locality condition, is not a model according to [11]. This demonstrates that the inheritance theory developed in [11] is nonlocal. Furthermore, this example illustrates that we need to know the context, in the form of inheritance paths, to determine the inheritance relation.

Another consequence of the lack of locality property of [11] is the counter-intuitive interpretation of certain network topologies [12]. For instance, suppose C's and D's are E's; B's are C's and D's; A's are B's and not D's. According to [11], there is an ambiguity about whether or not A's are E's; while the semantics in [2,6] support the conclusion that A's are E's, which is intuitively satisfactory.

[3] Later, we show, using Fig. 3, that the inheritance theory of [2] is intrinsically *nonlocal*.

[4] We also need to re-orient each arc to conform to the change in the direction of flow.

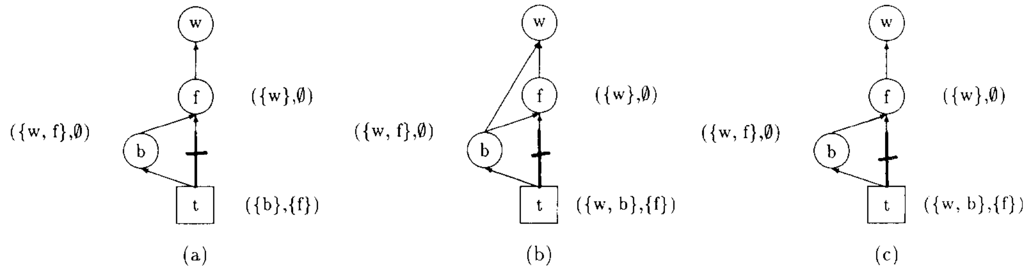

Fig. 2. Nonlocality of Touretzky's semantics II.

Note that our notion of locality is *semantic*, while that in [1,8] is *translational* [5]. These theories use meta-level axioms to determine specificity.

Now consider Fig. 3 from [3]. The *local specificity*, used to decide whether or not individuals a and b inherit property s (resp. $\neg s$) via node r (resp. node p), depends only on the inheritance relationship between the nodes p and r, and not on the individual inheritance paths from a and b to s. In Figs. 3(a) and 3(b), a p is an r. Hence, individuals a and b inherit $\neg s$ via p in both networks, according to local theories in [6]. In contrast, according to [2], the meanings assigned to s in Figs. 3(a) and 3(b), are different. In Fig. 3(a), a p is a q and a q is an r, and hence, p is more specific than r. So, the arcs from a and b to r can be viewed as "redundant". Thus, both a and b inherit $\neg s$ from p. On the other hand, in Fig. 3(b), a is not a q. So the arc from a to r can be interpreted as contributing independent evidence for a to be an r. In this situation, there is ambiguity about whether or not a is an s, because we have incomparable conflicting evidence for inheriting s through r and $\neg s$ through p. Thus, even though both a and b possess properties p and r, as far as inheriting s or $\neg s$ is concerned, p is more specific than r with respect to b, while p and r are unrelated with respect to a. In other words, in Fig. 3(b), the specificity with respect to node s for individuals a and b, are different. This motivates us to generalize the notion of locality to *ground locality* where the specificity is further parameterized by the "inheriting" node.

The *ground specificity relation* for node b with respect to node p (denoted $\prec_p^b$) is a binary asymmetric relation among the arcs of the node-environment of p. There exists a *ground isomorphism* between the node-environments Γ_p^1 and Γ_q^2 if there exists a bijection between the nodes of Γ_p^1 and Γ_q^2 such that the *adjacency relation* among the nodes, and for all nodes b, the *ground specificity relation for b* among the arcs of the node-environments, are preserved.

Definition 3. A semantics of inheritance networks is *ground local* if, for any pair of networks Γ^1 and Γ^2 containing ground isomorphic node-environments Γ_p^1 and Γ_q^2, whenever the meaning assigned to all the neighbors of node p in Γ_p^1 is identical to the meaning assigned to the "corresponding" neighbors of node q in Γ_q^2, it is the case that the assignment $\mathscr{M}_\Gamma^2[\mathscr{M}_\Gamma^2(q) \leftarrow \mathscr{M}_\Gamma^1(p)]$ satisfies Γ^2.

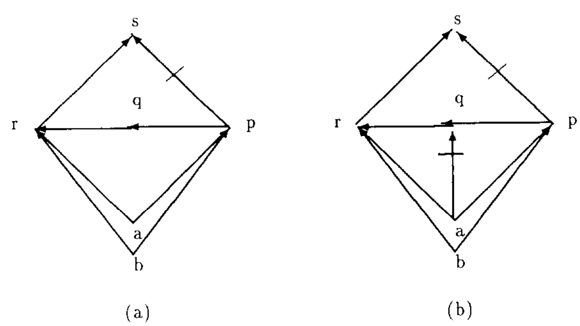

Fig. 3. Specificity can be nonlocal.

[5] That is, incremental changes to the network result in incremental changes to its translation.

Theorem 4. *The path based theory of* [2] *(resp.* [10]) *is* ground local.

Proof (Sketch). We derive an ordering on the in-arcs of a property node p, as a "function" of an "inheriting" node n, from the reformulated definition of *defeasible preemption* in [3].

A word about the notation used. $\pi(n, \tau_1, x, \tau_2, q)$ denotes a sequence of positive arcs from node n to node q via node x, and $\pi(n, \sigma, q) \rightarrow p$ (resp. $\pi(n, \sigma, q) \not\rightarrow p$) denotes a path from node n to node p ending with a positive (resp. negative) arc from node q to node p. Γ stands for an inheritance network and Φ for the extension constructed "so far" [3].

A positive (resp. negative) path $\pi(n, \sigma, q) \rightarrow p$ (resp. $\pi(n, \sigma, q) \not\rightarrow p$) is *preempted* in the context $\langle \Gamma, \Phi \rangle$ iff there is a node x such that (i) either $x = n$ or there is a path of the form $\pi(n, \tau_1, x, \tau_2, q) \in \Phi$, and (ii)$x \not\rightarrow p \in \Gamma$ (resp. $x \rightarrow p \in \Gamma$). Note that only the existence of the paths σ, τ_1 and τ_2 is required. So, we can translate these preemption criteria to ground local specificity as follows: $\langle q, p \rangle \prec_p^n \langle n, p \rangle$ and $\langle q, p \rangle \prec_p^n \langle x, p \rangle$. For each "inheriting" node n, these conditions impose an ordering on the in-arcs into node p.

The skeptical theory in [2] and the credulous theory in [3] can be obtained by plugging this specificity relation into the declarative framework developed in [5]. The equivalence results can be proved by induction on the *height* [6] of a node in the network. □

It is also interesting to note that our notions of local specificity and ground local specificity resemble Horty's notions of *general subsumption* and *off-path preemption* respectively, as described in [3].

However, the semantics of [11] is still *ground nonlocal* as it is impossible to define a satisfactory ordering on the in-arcs of p in Fig. 1(b), and on the "reversed" out-arcs of t in Fig. 2(a). In Fig. 1(b), the orderings $\prec_p^a$ and $\prec_p^b$ for the individuals a and b respectively, cannot be independent because of coupling; while in Fig. 2(a), a suitable ordering $\prec_t^a$ cannot exist.

[6] In [2], height of a node is referred to as its *degree*.

Locality accrues benefits such as compositionality. To understand the meaning of the whole network, one can understand the meaning of its constituent subnetworks and then compose their meaning appropriately using only the meaning of the "interface" nodes. That is, one does not have to recompute the meaning of the whole network *ab initio* [7]. This enables a declarative understanding of the network and supports efficient parallel computation of the meaning.

Locality has practical significance in various subareas of Knowledge Representation. Locality allows a "context-free" application of an inference rule. Given a Prolog rule and the fact that all the body literals in the rule are proven, the head of the rule can be asserted. Similarly, given the credibility of each MYCIN rule, the certainty factors associated with the body literals, and the combination function, the certainty of the head can be determined. In probabilistic reasoning, conditional probabilities are not amenable to such "local" treatment. However, belief propagation in causal networks requires only local computation [7]. The value of a belief parameter associated with a node depends only on the corresponding parameter values associated with its children/parent nodes and the local conditional probability matrix. This is similar to specifying the contexual "support" and "defeat" information in a preferential network by relating the meanings of a node and its neighbors. The belief propagation algorithm requires a number of passes through the belief network for convergence to a stable assignment. In contrast, the local semantics in [6] can be computed in a monotonic fashion in just two passes because arcs of the form $\neg p$ to q (resp. $\neg q$) are absent.

Acknowledgement

I wish to thank Michael Kifer, David S. Warren, Richmond Thomason and the anonymous referees for their helpful comments.

[7] We are tacitly assuming that the gluing process does not change the specificity relationships among the nodes in the subnetworks.

References

[1] M.L. Ginsberg, A local formalization of inheritance: Preliminary report, Unpublished manuscript, 1988.

[2] J. Horty, R. Thomason and D. Touretzky, A skeptical theory of inheritance in nonmonotonic semantic networks, *Artificial Intelligence* **42** (1990) 311–348.

[3] J. Horty, Some direct theories of nonmonotonic inheritance, in: D. Gabbay and C. Hogger, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming* (Oxford University Press, Oxford, 1992).

[4] K. Thirunarayan, An analysis of property flow view vs individual flow view of inheritance, in: *Proc. ISMIS-91* (1991) 256–265.

[5] T. Krishnaprasad, M. Kifer and D.S. Warren, On the declarative semantics of inheritance networks, in: *Proc. 11th IJCAI*, pp. 1093–1098, 1989.

[6] T. Krishnaprasad, The semantics of inheritance networks, Ph.D. Dissertation, SUNY at Stony Brook, 1989.

[7] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, Los Altos, CA, 1988).

[8] H. Przymusinska and M. Gelfond, Inheritance hierarchies and autoepistemic logic, in: *Proc. 6th ISMIS* (1989) 419–429.

[9] B. Selman and H.J. Levesque, The tractability of path-based inheritance, in: *Proc. 11th IJCAI* (1989) 1140–1145.

[10] L.A. Stein, Resolving ambiguity in nonmonotonic inheritance hierarchies, *Artificial Intelligence* **55** (1992) 259–310.

[11] D. Touretzky, *The Mathematics of Inheritance Systems* (Morgan Kaufmann, Los Altos, CA, 1986).

[12] D. Touretzky, J. Horty and R. Thomason, A clash of intuitions: The current state of nonmonotonic multiple inheritance systems, in: *Proc. 10th IJCAI (1987) 476–482.*