The Knowledge Engineering Review, Vol 9:1, 1994, 21–55

Research developments in multiple inheritance with exceptions

PETER W. EKLUND

Department of Computer Science, The University of Adelaide, Adelaide 5005, Australia

Abstract

The inheritance problem can be simply stated: for any instantiation of an inheritance network, say a specific hierarchy Γ, find a conclusion set for Γ. In other words, find out what is logically entailed by Γ. This can be done in two ways: either by defining a deductive or proof theoretic definition to determine what paths are entailed by a network; or by translating the individual links in the network to a more general nonmonotonic logic and using its model and proof theory to generate entailments that correspond to what one would expect from "viewing" the inheritance hierarchy. Two approaches to a solution to the inheritance problem structure this paper. The first is widely known as the "path-based" or "proof theoretic", and the second, the "model-based" or "model theoretic". The two approaches result in both a different interpretation of default links as well as a variation in the entailment strategy for a solution to the inheritance problem. In either case, the entailments produced need some intuitive interpretation, which can be either credulous or skeptical. The semantics of both skeptical and credulous inheritance reasoners are examined.

1 Introduction

More fundamental goal of inheritance research is to find a reasonably intuitive formal semantics for multiple inheritance hierarchies. By formal semantics we mean Tarskian semantics, or some variation of model theory where the relations between expressions in a language and objects in the world to which they correspond are described. The semantic theory provides the proper context in which to adjudicate differences of opinion when inheritance intuitions conflict. It is conceivable that different semantic theories may be appropriate for different uses of hierarchies.

To establish the terminology of this survey I examine existing nomenclatures for multiple inheritance in AI systems. This allows a categorization of the major approaches for multiple inheritance by examining the semantic interpretations given to "default" inheritance links. The standard principles of multiple inheritance are described, namely "specificity", "path preemption" and a new term, called "path weakness", is defined. The semantic difficulties which arise from using these definitions are examined.

The two solutions to the inheritance problem that structure this paper are the "path-based" and "model-based" approaches. These result in different interpretations of default links and a variation in the entailment strategy for the inheritance problem. In both cases, these entailments need intuitive interpretations, either "credulous" or "sceptical". The semantics of both sceptical and credulous inheritance reasoners are examined in this survey.

Logics for multiple inheritance are examined by first showing the inadequacy of a mapping between the semantics of multiple inheritance and first order logics. What follows is an examination of the inheritance problem in terms of each of the major schools of defeasible reasoning. A final summary includes a table covering each of the major treatments covered.

1.1 Justifying multiple inheritance

Major conceptual difficulties in inheritance result when tree-like taxonomies break down, permitting multiple inheritance hierarchies. This happens both in engineering domains as well as in systems based on natural sciences. As an example, consider the following context tree from MYCIN and the subsequent comments by Buchanan and Shortliffe (1984):

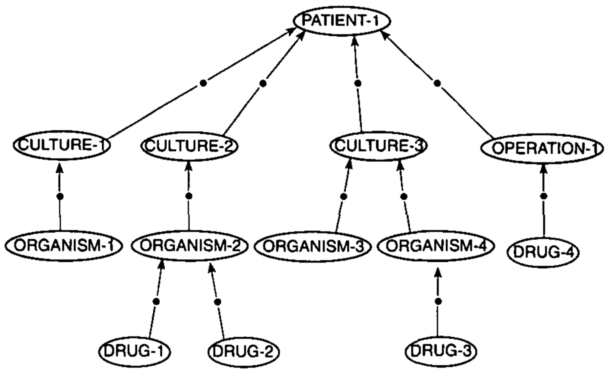

Figure 1 A fragment of MYCIN's context tree

"The context tree used by MYCIN is the source of one of the system's primary problems in attempting to simulate the consultation process. Every node in the context tree leads to the uppermost patient node by a single pathway. In reality, however, drugs, patients, organisms and cultures are not interrelated in this highly structured fashion. For example drugs are often given to cover for more than one organism. What we need, therefore, is a network of contexts in the form of a graph rather than a pure tree . . .". (Buchanan & Shortliffe, 1984, p. 132.)

It was in these early systems that the practical problems of multiple inheritance in knowledge engineering first arose as two or more taxonomic hierarchies combined into single knowledge-bases. One surprising result, which has yet to be widely known or accepted, was that first order logic was shown to be inadequate as a semantic basis for inheritance reasoning, even for relatively trivial inheritance hierarchies[1].

Further problems arose when exceptions to taxonomic hierarchies were incorporated into existing models. Exceptions arise in domains as diverse as database theory (Borgida & Williamson, 1985) and office information systems (Hägglund, 1989). Common to all such treatments is the idea of the progressive refinement of hierarchical schemas using exceptions.

Even in biology, where genus/species taxonomies exemplify taxonomic inheritance, exceptions occur as the boundaries of knowledge extend into the unfamiliar. For example, the discovery of the *Platypus* during European colonisation of Australia was received as a bad taxidermic joke when presented to the Royal Society in London. A warm-blooded animal, the Platypus has a duckbill and webbed feet, lays eggs and suckles its young. Once its authenticity had been verified, the Platypus cause a minor revolution in eighteenth century biological classification because it did not fit the existing taxonomy and represented a species belonging to a completely new *genus*, previously undiscovered, called the *Marsupials*. The introduction of the new *genus* into the existing taxonomy stabilised the theory and allowed hundreds of new species to be successfully classified. The moral of this story? If exceptions cannot be accommodated into existing taxonomies then they should force the review of the scientific theories on which the taxonomies are based. Exceptions thus motivate or encourage new theories. It could also be argued that the introduction of exceptions into a theory suggests the creation of a new class: as was the case of the Platypus and its *genus* the Marsupials.

[1]Section 5.1 discusses how and why this breakdown occurs.

Given the role of taxonomies in scientific theories, as an expression of the current state-of-the-art in a particular theory, it is not surprising that they occasionally require review. On the other hand, there are domains with naturally rigid taxonomic structures, medical domains for example, which are resistant to theory review[2]. This is due to historical or *lingua franca* factors beyond the influence of changes in the theories themselves[3]. Such rigid taxonomic structures are a justification for accommodating exceptions without enforcing theory review. In such cases, the existing theory should be preserved and augmented with defeasible inheritance links, thus accommodating both new and old theories alike. Such defeasible links represent normality, but may be defeated by exceptional cases.

1.2 What is an inheritance hierarchy?

An inheritance hierarchy covers a broad range of network knowledge representations which can be characterised as either directed acyclic graphs or trees. Their semantics are determined by the semantics of the arc types they contain, the meaning of the vertices in the graph, and the paths entailed from the network. The vertices within the hierarchy represent either individual objects or classes of concepts, which are connected by directed arcs defining subclass or superclass relationships. Legal paths, which can be entailed from an inheritance hierarchy, determine inheritable properties and are dependent upon: the semantics of the arcs connecting individual nodes; the semantics of the nodes themselves; and the intuitions which constitute an appropriate interpretation of the network.

Inheritance systems have given rise to a variety of nomenclatures for inheritance hierarchies. The following table summarises the terminology used for describing inheritance hierarchies and link types in various well-known systems. It fixes the terminology that will be used in this paper. As can be seen from Table 1, this treatment uses the term "inheritance hierarchy"[4] to refer to all the terminology types of inheritance hierarchy previously mentioned in column 4 of Table 1, and the "is-a" label to represent the inheritance links established between both class-to-class and instance-to-class links. Terms like "taxonomies" and "networks", which usually describe more general variations of inheritance hierarchies, will be avoided.

1.3 Normative statements

Although there can be little discussion over the meaning of the statement "p is always a q",[5] there are various interpretations of "normally p inherits the properties of q". Such "normative" statements are necessary if one intends to accommodate exceptions into multiple inheritance hierarchies. As a result, one of the major semantic issues in multiple inheritance is that of "normative" of "default" inheritance links. In other words, what interpretation does not give to the statement "normally p inherits the properties of q".

Touretzky (1986) makes the most comprehensive statement regarding the general intuitions of normative statements. Intuitively, the normative link $p \rightarrow q$ could be interpreted: typically p is a q; normally p is a q; nearly all p's are q's; p's can be assumed to be q's; and the default for p, w.r.t. q, is that p is a q. This linguistic mixture of normativity gives rise to two basic approaches to the inheritance problem, which in turn have three primary interpretations for normative statements.

[2]By "theory review" I mean the creation of new classes of objects or other actions designed to modify the inheritance structure.

[3]The taxonomy expresses objects and relations which provide a common framework for collective experience about the theory.

[4]The word "hierarchy" actually implies that nodes in the network inherit only once. In this sense, using "inheritance hierarchy" to refer to multiple inheritance is technically incorrect. It is, nevertheless, widely used, and we continue this convention.

[5]Such inheritance assertions are "strict" in the sense that they have a fixed interpretation, p and q can be interpreted as predicates which accept an individual a as an argument and whose conjunction is evaluated true, $p(a) \wedge q(q) = \mathrm{T}$.

Table 1 Nomenclatures of some inheritance systems

Language	*Researcher(s)*	*Link Terminology*	*Taxonomy Name*
SIMULA	Dahl, Nygaard	subclass/instance-of	subclass hierarchy
SMALLTALK	Kay, Goldberg	subclass/instance-of	class structures
LOOPS	Bobrow, Stefik	super	inheritance network
KL-ONE	Brachman	superc/individuates	taxonomy
			taxonomic structure
FLAVORS	Weinreb	component	tree of flavors
	Fahlman	ISA link	hierarchies
			tangled hierarchies
	Touretzky	instance assertion	inheritance hierarchy
NETL	Goodwin	supertype, subtype	inheritance hierarchy
			inheritance lattice
			taxonomic lattice
	Henrix	s/e	hierarchical taxonomy
	Winston	AKO/ISA	network
	Charniak, McDermott	ISA/inst	isa hierarchy
PSN	Levesque	ISA/instance-of	generalization hierarchy
			isa hierarchy
	Sowa	genus	type hierarchy
CMOL	Bubenko, Lindencrona	iss, ise, isa	isa hierarchy
	Smith	generalization	generalization hierarchy
RM/T	Codd	supertype	type hierarchy
	Quillian	subset	–
LINCKS	Padgham	is-a	inheritance
This treatment		*is-a*	*inheritance hierarchy*

1.3.1 Normality I

The first approach is that since normative statements have a multiplicity of intuitive interpretations, associating a formal semantic interpretation will result in a divergence from one or more of the intuitive interpretations. Therefore, no formal semantics for normative statements should be enforced. Furthermore, inheritance can be characterised without any explicit discussion of normative statements by examining the general intuitive semantics of inheritance independent of such interpretations. The results are usually fairly complicated deductive definitions of preferred inheritance paths via a global analysis of the network (Horty et al., 1990).

1.3.2 Normality II

The second approach is a local translation of normative links to a more general nonmonotonic logic by reinterpreting specificity as a statement of defeasibility. The model or proof theoretic semantics of a nonmonotonic logic can be used to infer preferred paths in the network which conform to the intuitive path entailments. Thus, normative statements are either directly encoded as axioms, as in Pearl (1988) and Doherty (1989), or as inference rules, the most obvious example being Etherington (1988).

In the second approach, local translations of normative links, opinion is divided about the formal interpretation for the normative statements between the probabilists and the rest of the community[6]. The first of these interpretations of normative links, and most vehemently argued against, is probabilistic. The idea is that a normative link $p \rightarrow q$ should have the interpretation: typically p is a q means that any randomly selected p is most likely to be a q, i.e. the conditional probability $P[q|p]$ is close to 1.

[6]For want of a better label this group can be described as the logicists (Pearl, 1988). This issue is not restricted to inheritance, but pertains to defeasible reasoning systems in general.

The counter argument to a probabilistic interpretation is that the implication that empirical and observable statistics are needed for the likelihood of events in the particular environment is too strong. For example, few would argue that in an equatorial region $P\,[Flies(x)|Bird(x)]$ is *high* but moving south to "Penguin Island" would necessitate a reduction in the probability of $P[Flies(x)\,|Bird(x)]$ as the frequency of encountered *penguins* increases. The probabilities associated with normativity are thus situation dependent. Furthermore, the logicists argue that the intended meaning of a defeasible link $p \rightarrow q$ is: in the absence of information to the contrary about some individual a other than $a \rightarrow p$ assume that $q(a)$ is true. Such assumptions can be made irrespective of the relative likelihood of their being true, and simply reflect an adherence to convention based on information at hand. By way of this definition, conventions as default rules are themselves domain dependent but adaptable to changing circumstances. These approaches appear in more detail later. In summary, there are three commonly held views regarding the meaning of normative statements:

- make no formal fixed interpretation of normative statements other than that $p \rightarrow q$ indicates that p is more specific than q. Treatments conforming to this idea are often called path-based approaches to inheritance reasoning:
- normative statements are treated as statements of fuzzily quantified or relative probabilities. Generally called probabilistic approaches to inheritance reasoning;
- normative statements are considered a set of axioms or inference rules in a nonmonotonic logic. These are often referred to as the model-based approaches to inheritance reasoning[7].

Generalising the above leads us to the following definitions of strict and default links in multiple inheritance. The definition of node types covers most treatments of multiple inheritance.

Definition 1 *An inheritance network* Γ *is a directed acyclic labelled graph* $\Gamma = (V,E)$. *The set of vertices* $V = \{v_1, v_2, \ldots, v_n\}$ *for* $n > 0$. *Any* $v \in V$ *has a type signature range* $\mathcal{R} = \{instance, class\}$ *and the set* $\mathcal{I} \subseteq V$ *is the set of individuals or instances with the type signature instance.* $\mathcal{I}$ *is labeled from the beginning of the alphabet, namely a,b,c,...* $\mathcal{C}$ *is a subset of all vertices* $\mathcal{C} \subseteq V$ *which have the type signature class.* $\mathcal{I} \cap \mathcal{C} = \emptyset$. *Classes are labeled from the middle alphabet p,q,r,s...*

The set of edges (E) in a multiple inheritance hierarchy $\Gamma = (V,E)$ *range over the link types* $\{\Rightarrow, \not\Rightarrow, \rightarrow, \not\rightarrow\}$ *with the corresponding labels* {*is-a*, *is-not-a*, *d-is-a*, *d-is-not-a*}. *These types are defined as follows*:

is-a strict inheritance: *If* $a \in \mathcal{I}$ *and* $p \in \mathcal{C}$ *then* $a \Rightarrow p$ *is read "a is a ground instance of p". This can be interpreted by considering p as a predicate in classical logic which accepts an individual a as an argument which evaluates to true,* $p(a) = \top$. *Graphically we can represent this as*:

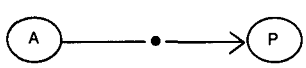

If $p,q \in \mathcal{C}$ *then* $p \Rightarrow q$ *is read "p's are always q's". This can interpreted if* $a \in \mathcal{I}$ *and* $p,q \in \mathcal{C}$ *and* $p(a) = \top$ *then* $q(a) = \top$. *This is drawn*:

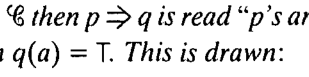

is-not-a strict class exclusion: *If* $a \in \mathcal{I}$ *then* $a \not\Rightarrow p$ *is read "a is not a ground instance of p". This can be interpreted by considering p as a predicate which accepts a as an argument which evaluates to false,* $p(a) = \bot$. *This is drawn*:

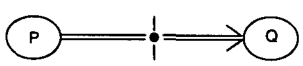

[7]It is worth pointing out that there is considerable inter-play between these three approaches, although the initial ideas on which they were based are different.

If $p,q \in \mathcal{C}$ then $p \not\Rightarrow q$ is read "p's are never q's". This can interpreted, if $a \in \mathcal{I}$ and $p,q \in \mathcal{C}$ and $p(a) = \top$ then $q(a) = \bot$. This is drawn:

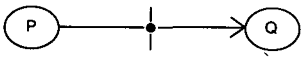

d-is-a default class inheritance : *$p \rightarrow q$ reads $p,q \in \mathcal{C}$ then p's are normally q's but there are exceptions, i.e., if $a \in \mathcal{I}$ then $\exists a\ p(a) = \top$ and not $q(a) = \top$. This will be drawn as:*

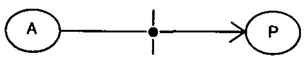

d-is-not-a default class exclusion: *$p \not\rightarrow q$ reads p's are normally not q's but there are exceptions, i.e., if $a \in \mathcal{I}$ then $\exists a\ p(a) = \top$ and not $q(a) = \bot$. This will be drawn:*

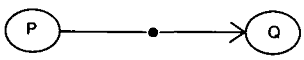

As a final word on normative statements, it should be mentioned that there is a significant school of thought that defeasible or default links should be kept out of any logical treatment of inheritance entirely. Proponents of this approach believe that normative reasoning should be incorporated into a process of belief revision. When one discovers that the original assumptions are wrong or inappropriate, then the theory's assumptions are revised and the entailment recomputed. This point of view is can be criticised on the basis of "cognitive adequacy" (Touretzky, 1984b), since it fails to differentiate between consequences drawn from default information and those drawn from facts known to be certain in the domain. The ability to distinguish the two may be vital in domains where we need to make decisions on the basis of reasoning with incomplete knowledge (Doherty, 1991).

1.4 Characteristics of inheritance hierarchies

Although there are rather obvious differences between acyclic digraphs and finite trees, it is worth pointing out that inheritance hierarchies described by them have radically diverse characteristics that are sometimes overlooked. For example, Thomason et al. (1987a) argue that the almost universal belief that strict inheritance has a mapping to first order logic is incorrect, without stating the underlying assumption; that he is specifically referring to multiple inheritance as the obvious counter example.

There are seven major attributes of inheritance hierarchies that can be used in combination, as summarised in Table 2.

- **Finite Tree Inheritance Hierarchies** can be described in graph theoretical terms as a finite tree. This means any one node has at most one parent, and there are a finite number of nodes. The implications are that there are no infinite paths in the hierarchy, and that there is a single path between any node pair, or no path at all. If a path is contained in the hierarchy, it is is entailed by the hierarchy. For example, Fig. 1 is such a network. Inheritance hierarchies described by finite trees have a semantics which is equivalent to a classical propositional interpretation.

Table 2. Combinations of allowable inheritance hierarchy types

	Strict	*Defeasible*	*Mixed*	*Bipolar*	*Cyclic*
Finite trees	√	×	×	√	×
Acyclic diagraphs	√	√	√	√	√

- **Bipolar Inheritance Hierarchies** include both positive and negative link types. Statements of the form "Fred is not a spy" (*Fred* $\not\rightarrow$ *spy*) and "Ivan is a spy" (*Ivan* $\rightarrow$ *spy*) can be expressed in such networks. Bipolar inheritance hierarchies may be either strict, defeasible or mixed strict and defeasible.
- **Multiple Inheritance Hierarchies** are also called tangled inheritance hierarchies (Fahlman, 1979). The analogy between first order logic and inheritance breaks down. Any one node in the network can have multiple parents, and thus inherit properties which may, in bipolar multiple inheritance hierarchies, conflict with one another. For example, $a \rightarrow p$ and $a \rightarrow q$ can simultaneously be expressed in such networks.
- **Homogenous Hierarchies** contain only a single link type. Thus, all links in the hierarchy are either all strict or all defeasible.
  - **Strict Inheritance Hierarchies** are homogeneous hierarchies in which only strict inheritance links are used. This permits statements of the form "elephants are mammals" (*elephants* $\Rightarrow$ *mammals*) where the intuitive meaning is that all elephants, ***without*** exception, are mammals. Any path within a strict inheritance hierarchy is entailed by the hierarchy.
  - **Defeasible Inheritance Hierarchies** are homogeneous hierarchies in which only defeasible inheritance links are used. Defeasible inheritance permits statements of the form "Most elephants are grey" (*elephants* $\rightarrow$ *grey*) where the initiative meaning is that there may be exceptions to the general rule. The semantics of defeasible inheritance networks are considerably more complicated than strict inheritance because they do not entail all the paths implicit in the hierarchy. Paths may be invalidated by more specific paths which preempt them, and contradicting paths may result in neither being entailed.
- **Mixed Strict & Defeasible Inheritance Hierarchies** Although relatively little progress has been made in mixing strict and defeasible inheritance links in networks (Horty & Thomasen, 1988; Gelfond & Przymusinska, 1990; Horty, 1990a,b; Padghom, 1989b), clearly this is a major direction of research in the area. In these networks, both strict and defeasible links can be used in some legal combination. For example, we may want to express the fact that most elephants are grey (*elephants* $\rightarrow$ *grey*) and that all elephants are mammals (*elephants* $\Rightarrow$ *mammals*) in the same hierarchy. Such networks are often called heterogeneous.
- **Acyclic Inheritance Hierarchies** Most, if not all, intuitive inheritance hierarchies are acyclic although some treatments consider cyclic inheritance networks (Geffner & Verma, 1989). Although monotonic inheritance could be cyclic, $p \rightarrow q$ and $q \rightarrow p$, it causes no conflict and is simply an argument for collapsing the two classes p and q into a single node. In cases where mutually exclusive nodes are connected by negative links, $p \rightarrow q$ and $q \not\rightarrow p$ simultaneously, the fact that an inheritance path cannot proceed beyond a negative link means that the semantics are precisely that of direct contradiction $p \rightarrow q$ and $p \not\rightarrow q$. In fact, Horty and Thomason (1988) make their negative links bidirectional, $p \not\Rightarrow q \equiv p \not\Leftrightarrow q$ to explicitly indicate this symmetry. Cycles which include more than two nodes are difficult to justify intuitively in an inheritance framework.

1.5 Notation

While discussing the various inheritance formalisms, we adapt the notation in Horty et al. (1990) and modify the presentation of each formalism as necessary.

Γ_i refers to a particular network, the subscript corresponding to the relevant network figure in this survey.

As discussed, the meaning of normative or defeasible links is dependent on the particular formalisation under discussion. What is true of all approaches is that the normative $p \rightarrow q$ can be interpreted "p is more specific than q". Specificity is the name given to a topological sort on Γ where the nodes which can be reached from any arbitrary node p have a lower *specificity* that p. The specificity of any node in Γ can be determined from a standard graph theoretic result in which all vertices of an n node acyclic directed graph can be labelled according to their outgoing arcs with $1, \ldots, k, k \leq n$.

Definition 2a—Node Specificity *Given that all $p \in V$ are labelled such that all vertices that can be reached from p via a directed path are given a label which is indexically less than p's*: (i) *p's are more specific than nodes which have indexically lower labels*; (ii) *if two nodes p and q have the same label they belong to the same specificity class.*

Definition 2b—Link Specificity *The vertices in Γ are partially ordered according to Definition 2a. Any edge e in Γ has a specificity corresponding to the specificity of the least specific node involved in the link. This is called the specificity class of e.*

Letters at the start of the Greek alphabet such as $\alpha, \beta, \gamma \ldots$ represent complete links such as the proposition $\alpha = a \Rightarrow p$ or a sentence as in $\beta = p \Rightarrow q$. The lower case Greek σ, ϕ, ψ are paths whose links may range over a sequence of links, depending on the system being discussed, and are typically constrained by the definition a legal path.

1.6 Legal paths in Γ

The definition of path legality is open to various interpretations. In Touretzky and Thomason (1988), paths of the form $p \rightarrow q \Rightarrow r$ are permitted. This type of path construction can be justified on the grounds of a local translation of inheritance. For example, one can imagine that instances could be attached at any level in the inheritance hierarchy as shown in Γ_1:

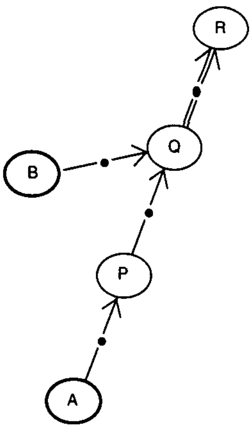

$= \Gamma_1$

The problem is that this leads to inheritance links which have both strict and defeasible interpretations depending on how the network is read. On the other hand, one can take the view that the interpretation of an inheritance link is singular and takes on the characteristics of the weakest interpretation of any subpath that it is involved in. In Γ_1, the link $q \Rightarrow r$ in the path $p \rightarrow q \Rightarrow r$, has exactly the same semantics as the defeasible $q \rightarrow r$. The argument is that if the inheritance reasoner is top-down, then q is strictly an r and defeasibility, p defaults to q, and thus p defaults to being an r, i.e., the impact of the strict link has been reduced to the weakest element in the chain of argument. Likewise, when reasoning bottom-up, p is by default and q but we cannot now say that p's are always r's[8]. On the other hand, the argument for allowing such a configuration is the network $\Gamma_1 = \{b \Rightarrow q, p \rightarrow q, q \Rightarrow r\}$, allowing $a \Rightarrow q$ and $q \Rightarrow r \supset a \Rightarrow r$ would be perfectly consistent with intuition. Contrary to this particular reading is that since the link $p \rightarrow q$ precedes $q \Rightarrow r$ we know that $q \Rightarrow r$ will, from time to time, be defeated and thus $q \Rightarrow r$ should be defeasible.

On the issue of negative links terminating an inheritance chain, Horty and Thomason (1988) give strict negative links the added semantic feature that they contrapose. This gives a certain symmetry

[8]If we were to reason in this way then $p \rightarrow q \wedge q \Rightarrow r \vdash p \Rightarrow r$, and symmetrically $q \Rightarrow r \wedge p \rightarrow q \vdash p \Rightarrow r$, which is clearly not the case.

to the intuitions behind the idea that one cannot inherit beyond a negative strict link. It is also an argument against cyclic inheritance networks. Almost all authors in the area accept that negative links (i.e., those of the form $\not\rightarrow$ and $\not\Rightarrow$) should terminate a valid chain of reasoning, except for Gelfond and Przymusinska (1990b). This implies that paths such as $p \Rightarrow q \not\Rightarrow r \Rightarrow s$ are impossible to construct, and are instead read as $p \Rightarrow q \not\Rightarrow r$.

Consequently, it seems reasonable to place two important restrictions on the semantics of multiple inheritance networks: (i) defeasible links of type $\{\rightarrow, \not\rightarrow\}$ cannot be followed by strict inheritance links of type $\{\Rightarrow, \not\Rightarrow\}$; (ii) negative links $\{\not\Rightarrow, \not\rightarrow\}$ can only appear in a legal path as the very last link, i.e., inheritance reasoning is terminated by one of these links.

A sequence of links is not a path if it contains a defeasible link followed by a strict link, or if it contains a negative link which is not the last link in the sequence. When the above conditions hold, a legal path between two nodes p and q can be written $\sigma_{p,q}$.

2 Principles for multiple inheritance

2.1 Consistency

The origin of the debate concerning the correct intuitive semantics for multiple inheritance with exceptions are so-called inconsistent networks, which are characterised by inconsistent states or contradictions. For example, if there is a pair $(p,q) \in \Gamma$ and two or more paths exist between p and q which have opposite polarity, Γ said to be inconsistent. Pearl (1987) offers a considerably more sophisticated definition of network inconsistency by considering whether a probability model exists for which all defeasible links are highly probable. If no such model exists, the network is considered to be inconsistent. For our purposes, the simple definition will suffice to illustrate the idea.

Consider the inconsistent network Γ_2, which is a homogeneous multiple inheritance hierarchy with exceptions. Γ_2 demonstrates the fact that a is inconsistent w.r.t. q. This is because there are two paths $a \rightarrow p \not\rightarrow q$ and $a \rightarrow q$ contained in Γ_2 which are in conflict.

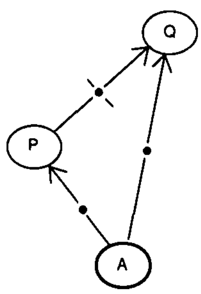

$= \Gamma_2$

The inconsistency is in itself no justification for ignoring Γ as a legitimate representation of the theory which Γ encodes. As it happens, inconsistency amongst paths can often be resolved, as it can for Γ_2, via the preemption criteria presented below.

2.2 Specificity

The criteria for determining network entailment is based entirely on the so-called "specificity principle". The idea is that one prefers more specific inheritance information than more general. In other words, an object inherits a property from the most specific inheritance class in the hierarchy. In bipolar multiple inheritance with exceptions, the specificity principle states that, when faced with a choice between competitive paths, arcs originating from more specific nodes have precedence over those from less specific nodes. Although specificity is a partial order on the vertices in Γ, it can only be used w.r.t. preferring Γ's paths. When a network contains multiple paths between a node pair, this criteria is used to establish a precedence relationship on those paths. Specificity thus produces a partial order ($<$) on nodes in Γ. To illustrate, in Γ_3 the specificity partial order is $p < q < r < s$.

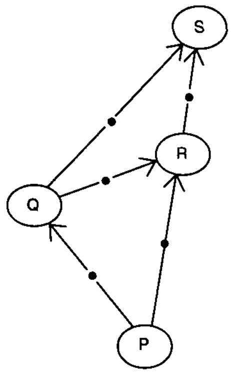

$= \Gamma_3$

2.3 On-path preemption

The concept of "preemption" results directly from the specificity principle. Preemption occurs when two competing paths rival for dominance by containing a node member which is more specific than the most specific node common to the two paths. A path σ is said to present a path ϕ if it contains a node which is more specific than all the other nodes that σ and ϕ do not have in common. Originally called preclusion by Touretzky (1986), preemption can also be thought of as "mutual neutralization". The idea is illustrated in Γ_4.

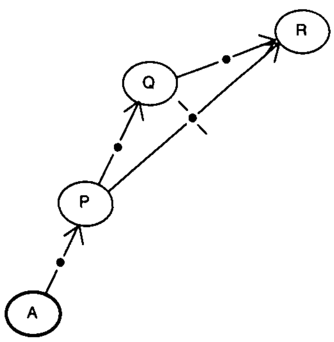

$= \Gamma_4$

In this example, the paths $\sigma = a \rightarrow p \rightarrow q \rightarrow r$ and $\phi = a \rightarrow p \not\rightarrow r$ are contradictory[9] and would neutralise each other so that nothing could be concluded. However, ϕ preempts σ, because $p < q$ and the reasoner should conclude ϕ. This is an example of "on-path" preemption where a path is preempted if one of its member nodes is involved in a counter argument. In this case, σ is rejected because one of its member nodes p is involved in a counterargument ϕ: ϕ subsequently preempts σ.

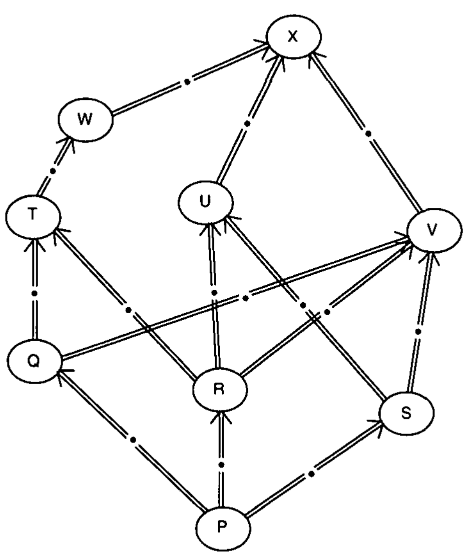

$= \Gamma_5$

[9]Contradiction or opposite polarity between competing paths is not ordinarily a precondition of applying preemption, although I argue it should be.

I argue that path preemption should only apply in inconsistent inheritance topologies. If preemption were to apply in consistent topologies, the dominant path in Γ_5 would be $p \rightarrow q \rightarrow v \rightarrow x$. If we then interpret the node labels as p = *square*, q = *rhombus*, r = *cyclic inscriptable quadrilateral*, s = *parallelogram*, t = *trapezoid*, u = *cyclic quadrilateral*, v = *inscriptable quadrilateral*, w = *quadrilateral*, then there appears to be no obvious reason why we should only conclude $p \rightarrow q \rightarrow v \rightarrow x$. The natural interpretation of Γ_5 is that a *square* (node p) has all the properties of the labels it inherits.

Consequently, providing we can find a mechanism which can detect inconsistent network topologies[10], there is no harm in pursuing a more liberal inheritance policy in consistent networks, and no need to apply path preemption when a network is consistent.

2.4 Off-path preemption

The distinction between on- and off-path preemption is illustrated in Γ_6. On-path preemption occurs when the path $\sigma = a \rightarrow p \rightarrow r \rightarrow s$ is preempted by the path $\phi = a \rightarrow p \not\rightarrow s$. Off-path preemption occurs when the same path ϕ preempts $a \rightarrow q \rightarrow r \rightarrow s$.

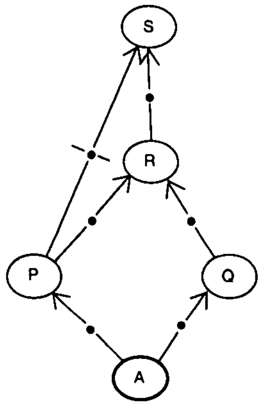

$= \Gamma_6$

Another idea necessitated by mixed strict and defeasible multiple inheritance is that of imposing a partial order on strict and defeasible links which belong to the same link specificity class. This introduces the idea of path "weakness". Weakness corresponds to the following idea, as illustrated by a modified version of a network due to Schlechta (1989), Γ_7.

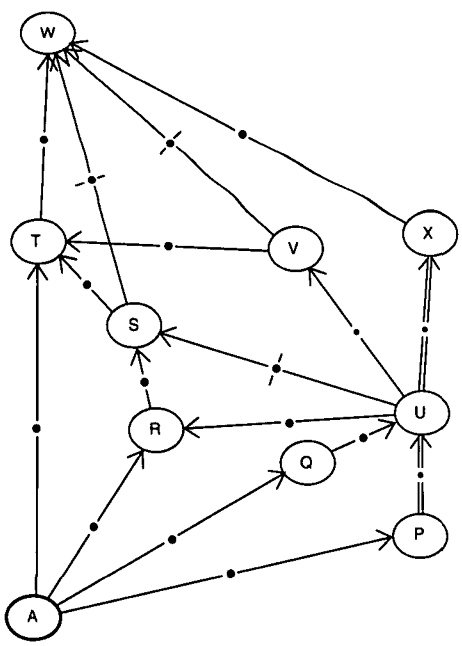

$= \Gamma_7$

[10]Detection of inconsistency in multiple inheritance is straight forward (Eklund, 1991).

Note that if we apply the usual preemptive criteria, the two credulous paths $\sigma = a \Rightarrow p \Rightarrow u \Rightarrow x \rightarrow w$ and $\phi = a \Rightarrow p \Rightarrow u \rightarrow v \not\rightarrow w$ are equally preferred. "Weakness" corresponds to the intuition that σ is a stronger argument for a is a w than ϕ is for a is not a w, and should thus be preferred. This corresponds to a preference relation between two links in the same specificity class. If one is defeasible and the other strict then the strict link is preferred.

This leads us to a point where we can now define network "entailment". Based on the above definitions, it is now possible to define a "support" relation for inheritance reasoning analogous to logical entailment. In a consistent network a sentence $\alpha = p \supset q$ is said to be supported by a network Γ if there is a legal path σ between p and q.

Combining the notions of support and path weakness[11] for bipolar mixed (strict and defeasible) multiple inheritance results in the following definition of path admissibility:

Definition 3—Bipolar Path Admissibility *Given an inconsistent network* Γ, *a path* σ *between p and q is admissible iff;*

(i) *there is no path* $\varphi_{p,q} \in \Gamma$ *which on-path preempts* σ;

(ii) *there is no path* $\varphi_{p,q} \in \Gamma$ *which off-path preempts* σ;

(iii) σ *is not "weaker" than any other admissible path between p and q.*

In other words, a path σ between two nodes p and q is said to be "admissible" if it is not preempted by any other path between p and q, and when competing credulous paths remain, σ is not "weaker" than any other credulous path. In such a case, σ is admissible and the p inherits the property q.

3 Multiple inheritance problems

3.1 Ambiguity

"Ambiguity" is a term coined by Touretzky (1984a) to describe a special case of inconsistency. Ambiguity describes an inheritance dilemma where compound conflicting paths in an inconsistent network necessitate an arbitary choice between competing inconsistent paths, originating from a particular node, because the preemption criteria fail to adjudicate the conflict.

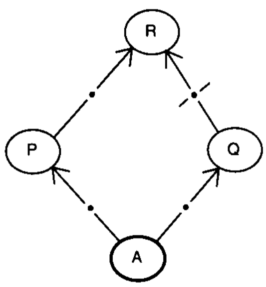

$= \Gamma_8$

To illustrate this idea, consider Γ_8, the topology of the so-called Nixon diamond[12]. This network would force an arbitrary choice whether an inheritance reasoner should choose one of $a \rightarrow p \rightarrow r$ or $a \rightarrow q \not\rightarrow r$, decide that both are true, or alternatively draw no conclusions given that the information is inconsistent. The network is thus said to be ambiguous, although more correctly the network is ambiguous w.r.t. the node a.

3.2 Updating inheritance hierarchies

An entailment procedure must not only draw the appropriate intuitive extensions from a given network but also behave consistently during update, i.e., monotonically $Cn(\Gamma) \subseteq Cn(\Gamma \cup \alpha)$. This is an appropriate juncture to discuss the origin of the often used name for multiple inheritance hierarchies, nonmonotonic multiple inheritance hierarchies. Update in inheritance hierarchies is

[11]Given two credulous paths σ, ϕ between a node pair with edges $\alpha = e^*_* \in \sigma$ and $\beta = e^*_* \in \phi$, σ is said to be "weaker" than ϕ iff; (i) α, β are in the same specificity class; (ii) α is defeasible and β is strict.

[12]The Nixon diamond (Reiter & Crisuolo, 1981) results from assigning the labels a := *Nixon*, p := *quaker*, q := *republican* and r := *pacifist* to Γ_8.

generally monotonic for instance update, i.e., if $\alpha = a \rightarrow p$ then $Cn(\Gamma) \subseteq Cn(\Gamma \cup \alpha)$, but nonmonotinic for class-to-class update, i.e., if $\alpha = p \rightarrow q$ then $Cn(\Gamma) \subseteq Cn(\Gamma \cup \alpha)$ may not hold.

Cross and Thomason (1987) use inheritance hierarchies as an example by which to investigate conditional logic as a framework for knowledge update. This logic is used to tell whether or not a given statement holds after the addition of an instance link to an inheritance hierarchy. Cross' work results from the "atomic stability theorem" (Horty et al., 1990). He treats instance link update in networks as the simplest form of update imaginable in an inheritance network, unable to upset the existing network entailment because for any instance link $\alpha = a \Rightarrow p$, if $\Gamma \vdash \alpha$ then for any statement β s.t. $\Gamma \vdash \beta$, $\Gamma \cup \{\alpha\} \vdash \beta$[13].

In Γ_9, $a \rightarrow p \rightarrow q \not\rightarrow s$ preempts $a \rightarrow p \rightarrow q \rightarrow r \rightarrow s$. Using a sceptical reasoner, the addition of $a \rightarrow r$ would have no effect on the entailment, $a \rightarrow p \rightarrow q \not\rightarrow s$ would still be the preferred path. The atomic stability theory tells us that despite instance link additions to the network the entailment will remain consistent and that instance link update is monotonic when using a skeptical reasoner.

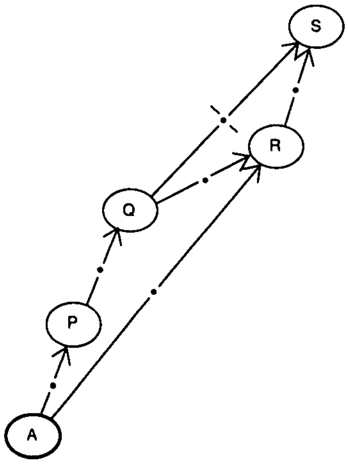

$= \Gamma_9$

The lack of generic stability in a sceptical reasoner is an argument against the general deductive view of inheritance. Any logic based on such inheritance semantics, in particular one which takes into account the semantics of redundant link addition, may not be semi-decidable. Despite this, the lack of general stability does not preclude semi-decidable logics with non-cumulative consequence relations.

3.2.1 Redundancy

"Redundancy" is the name given by Touretzky to the first of two semantic[14] problems with multiple inheritance systems.

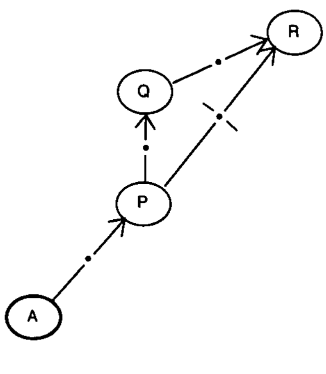

$= \Gamma_{10a}$

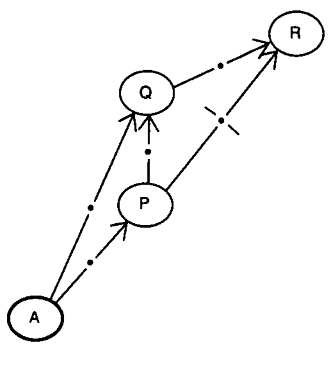

$= \Gamma_{10b}$

[13]This can be described as a lack of general "stability", or more correctly, non-cumulativity (Gabbay, 1985), if *Cn* produces a nonmonotonic consequence set, let $\alpha\ \beta \in Cn(\Gamma)$, then it is not necessarily the case that, $\beta \in Cn(\Gamma \cup \{\alpha\})$.

[14]There is an argument that this is a syntactic problem which relates to Touretsky's conditioning algorithm. However, the problem of redundant link addition is a demonstration of the general semantic problem of non-cumulativity which holds for multiple inheritance with exception.

If we consider the network Γ_{10a}, there appears to be a contradiction, there being a conflict between the inheritance paths $a \rightarrow p \rightarrow q \rightarrow r \rightarrow s$ and $a \rightarrow p \not\rightarrow r$. This conflict is easily resolved by the preemption criteria, and most reasoners would subsequently promote the path $a \rightarrow p \not\rightarrow r$. The redundancy problem arises when one adds the link $\alpha = a \rightarrow q$, as shown in Γ_{10b}. Adding α should not effect the conclusions which can be drawn from Γ_{10b}, since α was entailed from Γ_{10a} via the path $a \rightarrow p \rightarrow q$. The semantics of multiple inheritance systems thus represents a departure from the behaviour of traditional logics in that the introduction of a new axiom to the network, which is an element of the transitive closure of the theory, violates the initial consequences of the theory. Thus, if $Cn(\Gamma) \vdash \alpha, \beta$ then it is not necessarily the case that $Cn(\Gamma \cup \alpha) \vdash \beta$.

One way of dealing with the problem is to ban the addition of redundant links completely. As a general strategy for network construction this would be unwise.

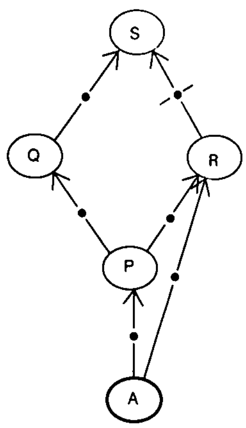

$= \Gamma_{11}$

In Γ_{11}, network ambiguity can be resolved through the addition of the link $\alpha = a \rightarrow r$ and the reasoner will entail $a \rightarrow r \not\rightarrow s$. If redundant links were forbidden, α could not be added to Γ_{11} due to existence of $a \rightarrow p \rightarrow r$. In some cases, the addition of redundant links should be encouraged rather than prohibited, particularly in ambiguous networks, as a simple and effective means of overcoming unnecessary nonmonotonicity. The problem with the argument is that there are two semantic interpretations of redundant links. The first is by way of the constructed inherited paths which the redundant link preempts, and the second is by way of the direct redundant link itself. Are redundant links simply extended structures resulting from Γ itself or are they actually representative of independently verified information? If they are artificially constructed from the existing network entailment, as in Touretzky's conditioning process, do we want them to override the existing entailment? The answer would be no. However, if redundant links represent independently verified information then they should be taken into account by the reasoner. Sandewall points out this problem referring to it as a "difference in derivational power" between links. He added that it is not merely a distinction between derived facts and axioms but of "derivational dependency", meaning the derivational power of a link is determined according to what other links are derivable from it.

To further illustrate the redundancy problem consider Γ_{12a}. $\Gamma_{12a} \not\Vdash a \rightarrow s$ since $a \rightarrow p \not\rightarrow r$ preempts $a \rightarrow p \rightarrow q \rightarrow r \rightarrow s$ and inheritance terminates at r. Since Γ_{12a} permits $q \rightarrow r \rightarrow s$ and so $q \rightarrow s$ the consequences of the addition of a direct link $q \rightarrow s$, as in Γ_{12b}, should be negligible. However, Γ_{12b} now permits $a \rightarrow s$ since $a \rightarrow p \rightarrow s$ preempts the link $a \rightarrow p \not\rightarrow r$. Horty describes such behaviour as, "like a situation in which the consequences of a set of axioms would be affected if the axioms were supplemented, not just with an arbitrary statement, but with a theorem derivable from those axioms". Such behaviour suggests that the graph-theoretic nature of inheritance reasoning provides an update semantics which are completely alien from deductive reasoning systems.

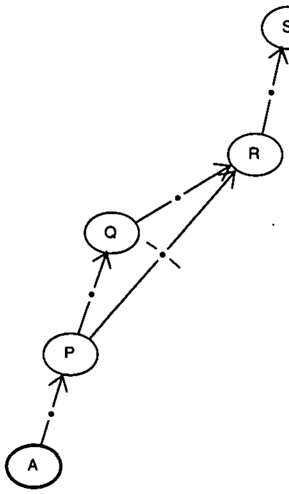

$= \Gamma_{12a}$

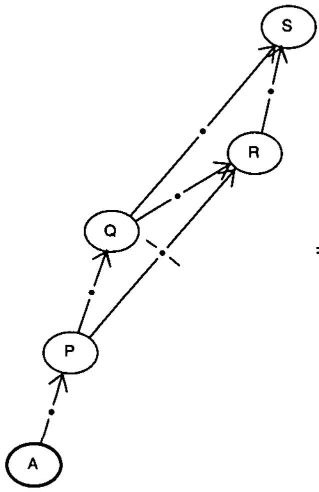

$= \Gamma_{12b}$

4 Path-based approaches

Having briefly examined the problems of multiple inheritance with exceptions, we now examine some of the intuitions which have resulted from research in path-based approaches to inheritance.

Thomason and Touretzky (1990) argue that the initial research efforts in developing algorithmic or path-based approaches to inheritance have been necessary to establish the semantics that we ultimately would like to capture in an axiomatisation of inheritance. The agenda is to "articulate the intuitions about reasoning with them [inheritance hierarchies] in a suitable form" (Thomason & Touretzky 1990). Once these intuitions are established, an appropriate model theoretic interpretation(s) need to comply with the technical criteria of adequacy, soundness and completeness. Finally, according to Thomason and Touretzky's timetable, the formalisation must demonstrate its practicality.

4.1 Shortest path reasoners

In a shortest path reasoner the response to a given query against a network is determined by a graph traversal algorithm which returns the shortest path from the source to the destination node. Consider Γ_{13}, for example; in response to the query "is a an s?" the reasoner would return the path $a \rightarrow p \rightarrow q \rightarrow r \rightarrow s$, the alternative path from $a \rightarrow s$ via $u \rightarrow v \rightarrow w \rightarrow x$ is ignored, being longer.

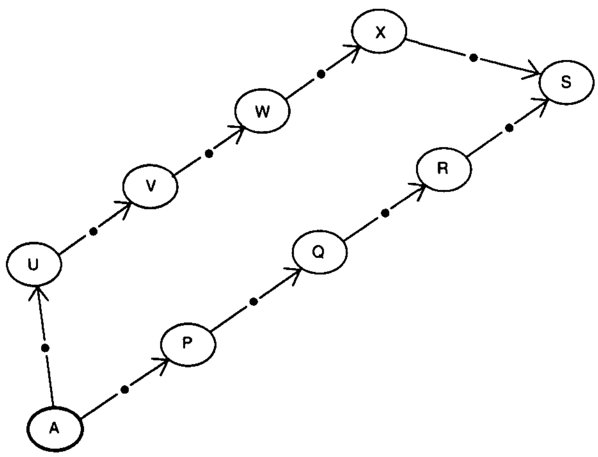

$= \Gamma_{13}$

According to Touretzky et al. (1987b), shortest path reasoners are provably correct for finite tree inheritance systems with or without exceptions, but fail in multiple inheritance networks in the presence of inconsistency and ambiguity. The deficiency of the shortest-path approach can be

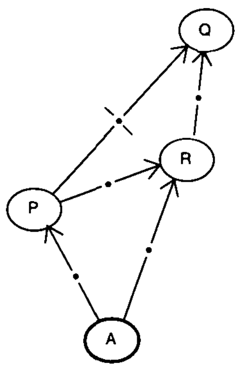

$= \Gamma_{14}$

illustrated by Γ_{14}, both $\varphi = a \rightarrow p \not\rightarrow q$ and $\phi = a \rightarrow r \rightarrow q$ are the shortest paths between a and q. However, φ preempts ϕ by way of $p < r$, and consequently φ should be entailed. This illustrates that shortest path reasoners violate the intuitions of the semantics of preemption.

In ambiguous networks, shortest-path reasoners will either report both inconsistent alternatives as being true or alternatively report the first extension it finds according to the search algorithm used.

4.2 Bottom-up versus top-down reasoners

The important thing about assembling the component paths from a network from direct links is that although preemption tell us whether a path in its entirety is disqualified, the reasoner must ensure that component sub-paths are not themselves ruled out.

Until Touretzky et al. (1987c), most approaches could be classified as being "top-down" or downward concatenation. Downward reasoners use as an inheritance metaphor "property flow" where an inheritance path is constructed from the general to the specific. The flow of properties can be interrupted by an exception. An inductive step would be equivalent to appending $p \rightarrow q$ to $q \rightarrow \sigma$ to form $p \rightarrow q \rightarrow \sigma$.

Alternatively, the "bottom-up" or upward concatenation metaphor is that of "argument construction", meaning that $\sigma \rightarrow p \rightarrow q$ is constructed by appending the direct link $p \rightarrow q$ to $\sigma \rightarrow p$. The most significant argument against downward reasoners is known as "decoupling", and occurs in sceptical reasoners. This phenomenon can be illustrated by Γ_{15}, which permits $a \rightarrow p \rightarrow q \rightarrow s$ since the subpath $a \rightarrow p \rightarrow r$ of the conflicting path $a \rightarrow p \rightarrow r \not\rightarrow s$ is preempted by $a \not\rightarrow r$. Since Γ_{15} permits $a \rightarrow p \rightarrow q \rightarrow s$, we can say that $a \rightarrow s$. Now consider the subpath $p \rightarrow q \rightarrow s$. Γ_{15} does not permit $p \rightarrow s$ since $p \rightarrow q \rightarrow s$ and $p \rightarrow r \not\rightarrow s$ conflict. Thus, the disallowed subpath $p \rightarrow q \rightarrow s$ is said to be "decoupled" from the path $a \rightarrow p \rightarrow q \rightarrow s$.

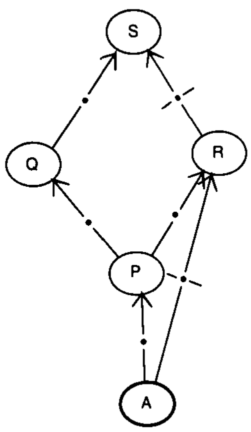

$= \Gamma_{15}$

According to the top-down approach, if $\Gamma_{15} \models a \rightarrow s$ then a must have inherited its "s'ness" from its immediate parent p. Since $\Gamma_{15} \not\models p \rightarrow s$, the "property flow" metaphor of the top-down approach fails. In addition to the problem of decoupling, downward inheritance reasoners are intractable (Selman & Levesque, 1989).

In the light of this evidence, one might wonder why downward reasoning is an option in the design space of an inheritance reasoner. The answer is aptly demonstrated by Touretzky and Thomason (1988). If we consider that other types of links, commonly called roles or relational links, extend inheritance into so called "terminological logics" (Patel-Schneidor, 1989; Bläsius et al., 1990; Nebel, 1990; Lenzerini et al., 1991)[15], real world problem domains can subsequently be represented as in Γ_{16}.

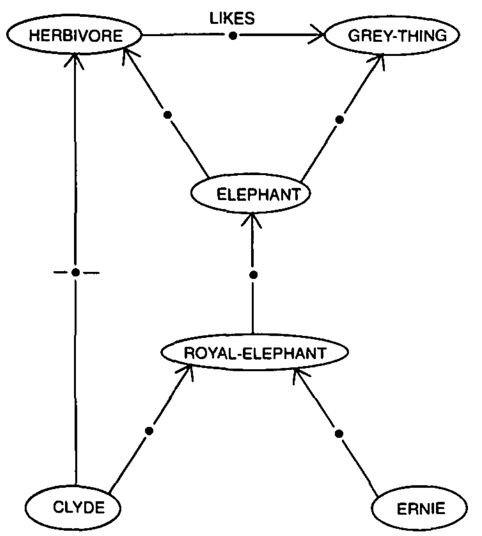

$= \Gamma_{16}$

Touretzky and Thomason (1988) introduce an extension of the standard inheritance vocabulary to express generic reflexive statements in the linguistic domain so that expressions like "royal elephants don't like themselves" are given a semantics in terms of inheritance hierarchy entailments. Irrespective of the application issues involved, one would want such a network to entail the path $c \rightarrow h \xrightarrow{R} g \leftarrow e \leftarrow r \leftarrow e$. To facilitate this type of construct, one would need both downward and upward path reasoners.

The upward approach does, however, solve the problem of decoupling, allowing properties to be coupled to their immediate parent nodes, and can also achieve polynomial performance (Selman & Levesque, 1989; Stein, 1990). For example, in Γ_{13} the path $a \rightarrow p \rightarrow q \rightarrow s$ cannot be formed by a downward reasoner, since it would require the formation of the subpath $p \rightarrow q \rightarrow s$. In this case, a coupling occurs between the properties that a can inherit and the properties that its immediate parent p can inherit.

4.3 Credulous inheritance

So far we have seen that a valid inheritance path in a network Γ is a combination of ideas. Firstly, it is an appropriate concatenation of paths using either a top-down or bottom-up reasoner. These paths should then be filtered so that paths which contradict one another and paths which are themselves preempted by other paths are excluded. We are then faced with the question as to what to do with competing or ambiguous paths. Should they cancel one another out entirely or should an inheritance reasoner consider each competing extension as an alternative explanation for Γ?

A "credulous model" (or opportunistic model) for inheritance reasoning is one where we consider all paths, or propositions represented by paths, which led from an instance a to a class p as supporting belief in the inheritance of the class p by the instance a. In the case of ambiguity, credulous inheritance reasoners resolve inconsistency in a network by constructing extensions which reflect different solutions to the resolution of conflicting paths.

[15]The likes of which describe a semantics for representation systems such as KL-ONE (Brachman, 1985).

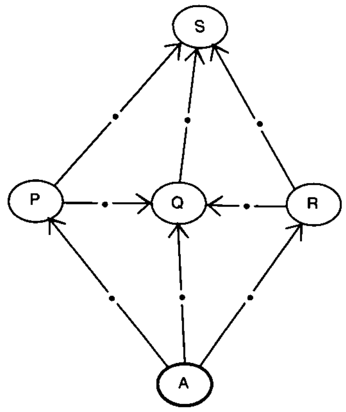

$= \Gamma_{17}$

In Γ_{17} we see that the assertion $a \rightarrow q$ is supported by three paths $a \rightarrow q$ itself, $a \rightarrow p \rightarrow q$ and $a \rightarrow r \rightarrow q$. The idea behind the credulous strategy is so called "belief hunger", where one tries to draw as many conclusions from a network as possible. At the same time that we can conclude that a is a q, we can also conclude that a has the properties of the classes p and r as well. When networks contain mutually exclusive paths, a number of consistent extensions are associated with paths through the network.

For an acyclic graph Γ, a legal path $\sigma_{p,q}$ between p and q is said to be credulous in Γ, written $\Gamma \triangleright \sigma$ iff $\sigma_{p,q}$ is a legal path and is not on- or off-path preempted.

4.4 Sceptical inheritance

The major objection to credulous inheritance is that it is computationally intractable, which leads us to consider alternative approaches which attach fewer models to inheritance hierarchies. The alternative to credulous models of inheritance are "sceptical" models, which encapsulate the idea that conflicting paths, representing conflicting arguments, cancel or neutralise one another. Sceptical inheritance describes the notion that a proposition p may be supported by a number of different lines of reasoning, but that only the facts present in all arguments for p should be used to defend it.

4.4.1 Blocking ambiguity

The definition of sceptical inheritance given by Touretzky et al. (1987c) is that a compound argument is neutralised by any conflicting argument which is not itself preempted. The definition allows the inheritance reasoner to rule out inconsistent compound paths like $a \rightarrow p \rightarrow q$ and $a \rightarrow p \not\rightarrow q$, but permits inconsistent direct links of the form $a \rightarrow p$ and $a \not\rightarrow p$[16]. The reason is the issue of deductive reasoning in the face of inconsistency.

The main aim of the sceptical inheritance research is to produce an inheritance reasoner which would encapsulate the idea of the intersection of all credulous extensions, called "ideal scepticism", without generating all credulous extensions to achieve this. Touretzky's et al. (1987c) first attempt at such a sceptical reasoner failed to give the same results as the intersection of all credulous extensions on the Nixon double diamond topology, and is an example of an "ambiguity blocking" sceptical reasoner. It is called ambiguity blocking (Stein, 1989) because as soon as it reaches an ambiguity the reasoner discontinues inheriting.

The general approach is illustrated on Γ_{18a}; if a node t is ambiguous w.r.t. a node p then all the arcs into and out of t are deleted. When the entire network has been scanned the remaining edges reveal a new network Γ_{18b}, which is unambiguous w.r.t. p and represents the ambiguity blocked sceptical extension of Γ_{18a}.

[16]The intuitions behind this result come from relevance logic, where inconsistent information cannot be used to derive arbitrary conclusions. In the case of direct contradiction, $a \rightarrow p$ and $a \not\rightarrow p$, the reasoner will be able to conclude $p(a)$ and $\neg p(a)$, but will not continue to draw irrelevant conclusions based on this inconsistency.

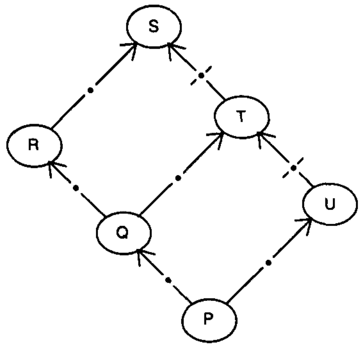

$= \Gamma_{18a}$

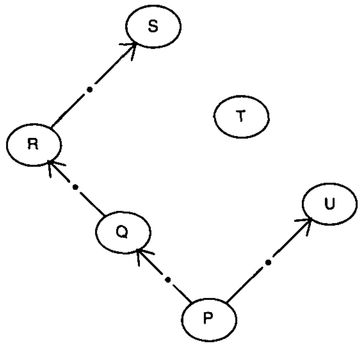

$= \Gamma_{18b}$

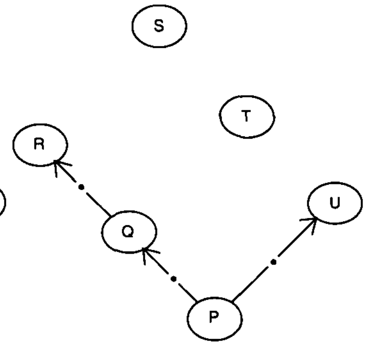

$= \Gamma_{18c}$

The trouble is that $q \longrightarrow t$ has been eliminated, which means that $\Gamma_{18b} \models p \longrightarrow s$, although it is not really intuitive to do so. If $q \longrightarrow t$ has been considered as an "in" link, then $\Gamma_{18b} \not\models p \longrightarrow s$. If we are to take ideal scepticism as the correct specification of a sceptical reasoning algorithm, then the ambiguity blocking algorithm fails the test.

4.4.2 Propagating ambiguity

The alternative is to propagate ambiguity, not by deleting all the incoming and outgoing arcs of an ambiguous node, but by marking such nodes ambiguous and continuing the network scan. Once the reasoner reaches the top of the hierarchy it backtracks over the network deleting all the arcs associated with ambiguous nodes it has marked.

The results are shown in Γ_{18c}, which represents the ambiguity propagated sceptical extension of Γ_{18a}. Although this approach appears promising for cascading diamond topologies, it fails for the following topology, whereas ambiguity blocking succeeds.

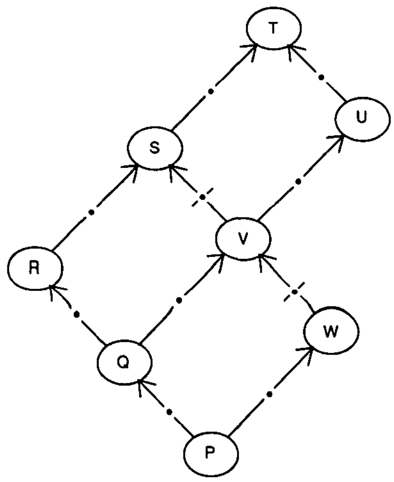

$= \Gamma_{19}$

The argument here is that in every credulous extension that one could consider $\Gamma_{19} \models p \longrightarrow t$[17], but ambiguity propagation rules this out. The conclusion is that there are certain network topologies in which a fact will be true in every credulous extension but have no justification in the intersection of these credulous extensions. Makinson and Schlechta (1989, 1991) call such facts "floating conclusions". Subsequently, any path-based approach to generating sceptical extensions will either

[17]Stein's labelling scheme is p = *seedless grape vine*, q = *grape vine*, r = *vine*, s = *arbor plant*, t = *plant*, u = *tree*, v = *fruit plant*, w = *fertile thing*, p = *seedless grape vine* is always a t = *plant*.

be unsound, in the sense that floating conclusions may be entailed with different arguments supporting each extension, or incomplete, in the sense that the floating conclusion is not entailed.

Subsequent treatments (Geffner, 1989) introduce alternative sceptical reasoners which produce the appropriate results from the double diamond topology of Γ_{18a}, but without any proof that it will satisfy all topologies.

Regardless of the debate between sceptical versus credulous inheritance reasoners, much effort has been expended to define a declarative mapping between a network and a logic representation which, given the appropriate axiomatisation and model semantics, can perform both sceptical and credulous reasoning within the same logical representation. After a brief historical summary of path-based approaches to multiple inheritance with exceptions, the following section examines attempts to achieve this goal via direct translations of inheritance networks into defeasible logics.

4.5 Path-based inheritance—Summary

The story of nonmonotonic inheritance or multiple inheritance with exceptions begins with work by Fahlman (1979) on a "parallel marker propagation" architecture called NETL. Inheritance networks could be represented by assigning a separate hardware element to each node and link. Every element consisted of connections to other nodes and links, and a small amount of memory by which markers could be passed on or propagated through the network.

The result was a fast and efficient means of storing and retrieving data stored in hierarchies. A problem arose when considering different types of inheritance hierarchy configurations. The NETL architecture and algorithms were well suited to finite tree inheritance structures, where any one node has at most a single parent in the hierarchy, but would fail when the hierarchy was extended to bipolar multiple inheritance systems. Fahlman's algorithms were unsound in the presence of ambiguous or redundant network topologies.

Touretzky (1984b) suggested that conflicting inheritance paths could either be resolved by way of a partial ordering criteria known as "inferential distance ordering"[18], or represented by way of "credulous paths", each of which expressed alternative and equally valid interpretation of the hierarchy in question. His solution was to condition the network, through the addition of links already entailed by it, via inferential distance ordering, which could disambiguate conflicting entailments. In doing so, the topology of the network would be modified and marker propagation algorithms would function properly.

Despite the soundness of parallel marker propagation algorithms after conditioning, reservations about the complexity of parallel marker propagation algorithms remained, particularly as they applied to NETL. As we have seen in section 3, there are also serious problems concerning the addition of redundant links during the conditioning phase and how they can modify the intended semantics of the network in a detrimental fashion.

Touretzky's work was subsequently published in book form (Touretzky, 1986), and this popularisation prompted a debate within AI circles as to what the "intuitive" and correct semantics of multiple inheritance with exceptions should be. Sandewall (1986) pointed out an inconsistency in the intuitions on which Touretzky's inheritance reasoner were based, namely Touretzky's omission of the possibility of off-path preemption, and thought that the semantics of inheritance, at least in the initial stages of the investigation, should be defined by a series of standard examples for which agreement about what should be entailed existed. He went further to produce his own specification for a credulous reasoner by way of a deductive definition of inheritance rules which could deal, among other things, with both on- and off-path preemption.

Not surprisingly, most inheritance research has not addressed defeasible reasoning in general, but concentrated on arguments based on intuitions involving the validity of paths in hierarchies and the inheritance reasoners which would conform to such intuitions. After a protracted and public discussion over four years, the definitive statement about these intuitions was made in Touretzky,

[18]The equivalent of a topological sort by specificity.

Horty and Thomason's (1987b) "Clash of Intuitions" paper. Here a design space for inheritance reasoning was defined. The issues of scepticism versus credulity, upward versus downward reasoners, the problem of decoupling in credulous inheritance, and on- versus off-path preemption were examined in detail.

It was commonly held up until this time that a sceptical reasoner would produce as an extension the result of the intersection of the credulous extensions. Makinson and Schlechta's (1989) study reached the conclusion that directly sceptical[19] approaches could not hope to replace credulous reasoners working with complete families of network extensions. The use of upward chaining sceptical inheritance reasoners leads to a violation of the specificity principle and an inability to enforce preemption.

Makinson and Schlechta go on further to challenge both the semantics and proof theory of sceptical inheritance. Two new terms were subsequently introduced into the literature, "floating conclusions" are those in which every acceptable extension supports a proposition but there is no common path or line of reasoning for the proposition. The other term is "zombie paths", where a sceptical inheritance reasoner is unable to eliminate a sceptically plausible path on the basis of the elimination of its subpath. In other words, sceptically acceptable paths may be supported by subpaths which are sceptically unacceptable, and this challenges the generality of sceptical inheritance reasoners.

If upward chaining sceptical inheritance reasoners are semantically incorrect, as Makinson and Schlechta maintain, then the future for sceptical inheritance is not particularly bright. Selman and Levesque's (1989) complexity results significantly reduce the design space of inheritance reasoners by showing that the only tractable reasoners will be those based on upward chaining inheritance. The combination of these results would seem to rule out a computable and semantically sound sceptical inheritance reasoner, although it is possible to imagine a heuristically driven downward reasoner whose worst case behaviour is intractable rather than *NP*-hard.

On a more positive note, since most of the negative side-effects of sceptical inheritance reasoning are well-known, namely floating conclusions and zombie paths, a user willing to accept such a system at least knows what the pitfalls are. This can be seen as a qualitative improvement over what was previously the case.

One of the motivations for a path-based semantics is that current model theoretic treatments of logical consequence in inheritance have yet to adequately characterise multiple inheritance, as we shall see in the next section.

5 Logics for multiple inheritance

Although Horty's et al. (1990) semantics of specificity, preemption and sceptical reasoning[20] conform to the expected intuitions for a large number of empirical cases, there are nevertheless no definitive semantics for all of the inheritance types covered in Table 2. As Horty et al. (1990) state, "the situation [in inheritance hierarchy semantics] is reminiscent of the situation in philosophical logic, where there exists rival logics embodying distinct conceptions of correct deductive reasoning." Under these circumstances, no single semantic formalism will suffice in every inheritance model, and a correct interpretation will be dependent upon the chosen inheritance semantics.

In the rest of this survey, I examine the intuitive semantics of inheritance hierarchies and related attempts to formalise the resulting intuitions. I begin by examining the possible mapping between certain monotonic inheritance hierarchies and standard logics demonstrating the difficulties which are involved. A probabilistic interpretation of inheritance gives a normative link $p \rightarrow q$ the interpretation that any randomly selected p is most likely to be a q, i.e., $P[q|p]$ is high. Treatments of multiple inheritance which deal with this probabilistic interpretation are examined in section 5.2.

[19]Path-based, algorithmic or deductive definitions.

[20]As we have seen in the previous section, there is still considerable discussion about what constitutes a semantically correct sceptical reasoner.

Following this, approaches which rely on a local translation of normative links to a more general nonmonotonic logic by reinterpreting specificity as a statement of defeasibility are examined. The model or proof theoretic semantics of a nonmonotonic logic can be used to infer preferred paths in the network which conform to the intuitive path entailments. Normative statements are either directly encoded as axioms, as in McCarthy and Hayes (1981), or as inference rules, as in Etherington and Reiter (1983). Default logic, conditional logic, autoepistemic logic, circumscription and preferential model treatments of multiple inheritance with exceptions are examined.

5.1 Standard logics and inheritance

One of the motivations for structuring knowledge in a hierarchical form is that it simplifies proof or inference procedures. The analogy between inheritance hierarchies and logical theories is that the total network is considered as a set of axioms or a theory, the nodes being propositional or predicate symbols, the links connecting nodes to be derivable consequences and the paths permitted in the network are like proofs. The idea is that graph traversal algorithms, whose behaviour is well known to computer scientists, can approximate generally intractable abstract proof theories by efficiently computing subclasses of the proof theories. The assumption is that the structural change in the knowledge, from say a propositional or predicate form, to a hierarchical representation increases the performance and efficiency of the proof procedure while leaving the original semantics unchanged.

As it turned out, a number of side-effects of this restructuring became apparent, and it was not always correct or intuitive to say that traversing paths in the network was equivalent to performing logical deductions in predicate calculus. Despite this, many authors (Etherington, 1988; Touretzky, 1984b; Boutilier, 1989), particularly those using translations of inheritance networks to nonmonotonic logic, have treated inheritance homogeneous inheritance networks as having a natural correspondence with first-order logic, citing early work by Schubert (1976) and Hayes (1977) in support.

Quite correctly, Thomason et al. (1987a) maintain that the belief that the logic of inheritance hierarchies is equivalent to first order predicate calculus is somewhat of a "folklore theorem" in AI. Some clarification is required here. Classical, finite tree, strict inheritance, which has no exception, is in fact equivalent to predicate calculus.

On the other hand, this result does not apply in multiple inheritance systems, those that Fahlman calls "tangled hierarchies". Multiple inheritance cannot be translated into either predicate or propositional logic. This result can be illustrated by considering the following mapping from homogeneous networks and monadic predicate logic;

$a \Rightarrow p \equiv p(a)$
$a \not\Rightarrow p \equiv \neg p(a)$
$p \Rightarrow q \equiv \forall x\, [p(x) \rightarrow q(x)]$
$p \not\Rightarrow q \equiv \forall x\, [p(x) \rightarrow \neg q(x)]$

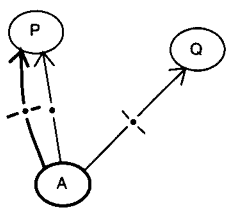

$= \Gamma_{20}$

Γ_{20} leads to the argument $p(a), \neg p(a) \vdash q(a)$, which means that whenever the contradictory form $a \Rightarrow p$, $a \not\Rightarrow p$ is present in the network, then *ex falso quodlibet* can be used to support statements which contradict sentences in the initial axiom set.

To labour the point, even if *ex falso quodlibet* does not apply, as in para-consistent relevance logics (Belnap, 1976, 1977), *modus ponens* presents problems, as illustrated by Γ_{21}. If this mapping were to hold then Γ_{21}, as a legal inheritance network[21], would entail $p(a), \forall x\, [p(x) \rightarrow q(x)] \vdash q(a)$ by *modus ponens*, which clearly contradicts $a \not\Rightarrow q$ in Γ_{21}.

[21]There is, however, a strong argument that Γ_{21} represents a semantically incorrect inheritance hierarchy. According to intuition, if it is homogeneous strict then inconsistency, as explicit contradiction, should be forbidden.

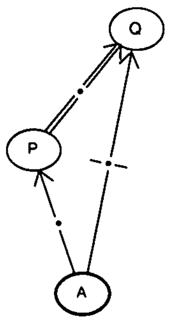

$= \Gamma_{21}$

This is, of course, the behaviour of an unsound theory w.r.t. an intuitive semantic theory of multiple inheritance. The result is that the above mapping between first order predicate calculus and inheritance hierarchies inadequately captures our intuitions concerning multiple inheritance. Interestingly, this example also provides us with evidence that there is no intuitive mapping between some classes of paraconsistent relevance logic and inheritance hierarchies. The conclusion is that even a logical characterisation of monotonic inheritance is considerably more complex than is widely believed, and that the semantics of classical predicate logic does not suffice as the semantics for monotonic multiple inheritance.

5.20 Probabilistic treatments

As mentioned in section 1.3, probabilistic treatments of multiple inheritance interpret the semantics of normative statements as fuzzily quantified or relative probabilities.

Bacchus (1988) develops his own probabilistic extension to first order logic, where the model structure includes a probability function over the domain of discourse and whose syntax is extended to allow the formalisation of terms which reflect statements of empirical probability. Like many formalists, he initially uses inheritance as an application by which to test the generality of the logic (Bacchus, 1988), and subsequently goes on to suggest his own inheritance reasoner conforming to the logic (Bacchus, 1989). Preference criteria by which to promote inheritance extensions take the form of universally quantified statements or strict inheritance links in the hierarchy. For example, in the "clyde elephant skip"[22], the hierarchy is encoded as follows:

$[Gray(x)|Elephant(x)] > c$[23]
$\forall x\ Royal\_Elephant(x) \rightarrow Elephant(x)$,
$[Gray(x)|Royal\_Elephant(x)] < 1 - c$,
$Royal\_Elephant(clyde)$, $Elephant(clyde)$.

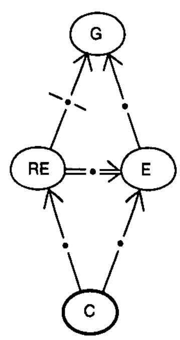

$= \Gamma_{22}$

On the basis of $\forall x\ Royal\_Elephant(x) \rightarrow Elephant(x)$ the preferred interpretation becomes a statement of belief,

$$B\ [Gray(clyde)|Royal\_Elephant(clyde)]$$

and this reflects the comparative belief in the support for the statement and cannot be probabilistically quantified. This turns out to be a problem for this logic in belief applications. In other words, the mapping between the logic and heterogeneous inheritance hierarchies is that defeasible links

[22]With the labels *G=gray*, *RE=royal elephant*, *E=elephant* and *C=Clyde*.

[23]*c* is an observable or empirical constant close to 1.

become statements of conditional probabilities and strict links are universal quantifications. This means that such a system does not sanction inheritance through more than a single defeasible link, a somewhat unituitive explanation of inheritance reasoning.

The main question, and the standard criticism resulting from Bacchus' logic, is where do we obtain the empirical probabilistic functions in the first place? Although Bacchus criticizes Pearl's logic, we should be sympathetic to the fact that Pearl's "system-Z" semantics are at least independent of such empirical probabilistic observations.

The general idea behind Pearl's treatment is that any set of default rules contains a natural ordering or prioritorization called a "z-ordering". All default rules are given a probabilistic translation, for example $P[q|p] \geqslant 1 - \varepsilon$ is the translation of $p \rightarrow q$.

In order for the z-ordering to be possible, a network should be globally consistent which means that an assignment can be found for all defaults such that $\varepsilon > 0$ but close to 0. A preferred model can then be found by considering the highest rank of the default rule whose ε value can be adjusted to falsify the model. This results in a preferred model which reflects an assignment to the default rules which is unique and can be calculated in polynomial time.

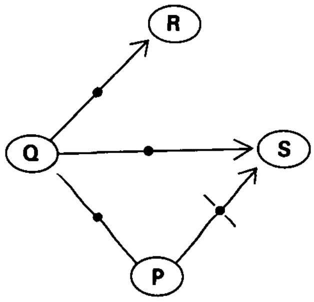

$= \Gamma_{23}$

The major weakness of this approach is the blocking of property inheritance across exceptional classes, for example Γ_{23} which Pearl argues should sanction p's and r's, on the grounds that $p \rightarrow q$ even though the p's are exceptional q's. To address this problem Pearl introduces an additional feature called "maximum entropy", where additional defaults are added to the preferred extension providing it remains consistent and that they increase the models entropy. In addition, the *z-rating* idea can be adopted to produce a partial order on models, in the style of Shoham, called conditional entailment to combat cross property inheritance.

5.3 Default logic and inheritance

In default logic a "default rule" of the form "if x, and it is consistent to believe y, then assume y" treats defaults as deductive inference rules external to the object language of the logic. Reiter's (1980) original proposal for default logic, and his subsequent treatment of multiple inheritance as an example application, was based on the assumption that default links could be represented by so called normal default rules of the form;

$$\frac{\alpha(x):\beta(x)}{\beta(x)}$$

Normal default theories have excellent computational properties, they guarantee that every normal theory has an extension, and are "semi-monotonic", i.e., $Th(d_1,\Gamma) \subseteq Th(d_1 \cup d_2, \Gamma)$[24]. This has the effect that any existing entailment of a normal default theory can be supplemented directly with new normal defaults without having to retract any conclusions previously made. The problem with normal defaults is that they may interact to produce anomalies (Reites & Crisuolo, 1981), e.g.,

[24]d_1 and d_2 are default rules and Γ is a set of first order axioms.

$$\frac{\alpha(x):\beta(x)}{\beta(x)} \quad \frac{\beta(x):\gamma(x)}{\gamma(x)}.$$

The problem in general is that default logic provides no way of distinguishing between a conclusion drawn from first-order axioms and those drawn from default rules. If we know $\alpha(a)$ then the first of the above rules concludes $\beta(a)$. This default consequence can subsequently be used in the prerequisite of the second rule to conclude $\gamma(a)$. The point is that $\beta(a)$ drawn from a default rule has a different intuitive semantics from that used in the prerequisite of the second default rule. The $\beta(x)$ in the first rule serves as a default conclusion, whereas in the second rule it is a fact.

This representational inadequacy can be partially overcome by introducing so called "semi-normal" defaults of the form

$$\frac{\alpha(x):\beta(x) \wedge \gamma(x)}{\gamma(x)}.$$

However, despite the fact that semi-normal default theories deal with the problem of transitivity among defaults, they have none of the desirable properties of normal default theories (Reiter & Crisuolo, 1981). Semi-normal default theories neither guarantee an extension nor are they semi-monotonic. The problem of finding extensions to semi-normal default theories is also shown by Selman and Kautz (1988) to be *NP*-hard.

Etherington and Reiter (1983) and Etherington (1987a) propose a semi-normal default theory for multiple inheritance via the introduction of a "cancellation" or "exception" link which, in itself has no independent semantics, but serves as an explicit mechanism for imposing a priority on default links[25]. In later work, Etherington (1987b) imposes an implicit ordering on default links using Touretzky's (1986) "inferential distance ordering". In so doing, he eliminates explicit cancellation links. Through the implicit ordering, which results from an examination of the prerequisite and consequence conditions of normal defaults, a partial order on defaults is defined. This results in a criteria for selecting an appropriate extension in the presence of ambiguity, and a return to the more desirable class of normal default theories. It also results in considerably more complicated proof procedure, and a prior global or non-local analysis of the normal defaults to establish the implicit ordering under inference.

In general, default logic treatments of inheritance suffer from:

- an inability to distinguish between an ordinary and a default conclusion;
- no obviously apparent translation from linguistic defaults to the formal representation, i.e. a choice of normal-, semi- or nonnormal defaults;
- the general class of default theories are intractable in the first-order case (Reiter, 1980): this also includes semi-normal default theories (Etherington, 1988).

5.4 *Conditional logic treatments*

Like Etherington, Boutiler (1989) argues that a formal semantic characterisation of inheritance can more readily capture intuitions and act as a yardstick by which to measure the performance of algorithmic techniques. Syntactically, links in a network are translated into logical statements in a conditional logic E in the obvious way;

$$p \rightarrow q \equiv p \supset q$$
$$p \not\rightarrow q \equiv p \supset \neg q$$
$$p \rightarrow q \wedge p \not\rightarrow q \equiv \neg(p \supset q) \wedge \neg(p \supset \neg q)^{26}$$

Boutiler bases his semantics on Shoham's (1988) preferential logic, and imposes a preference relation to define preferred models of the theory or network, Γ. The preferred models relation $\leq$ is defined in terms of the minimisation of contradictions constructed by considering alternate paths

[25]Froidevaux and Kayser (1988) offer a somewhat more sophisticated proposal along the same lines.
[26]Note that $\supset$ is conditional and not material implication.

and the specificity relation. Thus a preferred model is that which contains the most specific non-contradictory paths in Γ. The approach yields a sceptical, on-path preemptive reasoner which is stable, but does not cope appropriately with the addition of redundant links.

In his defence, Boutiler quite rightly points out that the way a reasoner should deal with redundant links depends on the intuitive semantics that they are assigned.

5.5 Autoepistemic logic and inheritance

The basic idea behind autoepistemic logic is that a nonmonotonic axiom can be written in the form $\forall x Bird(x) \wedge L(Canfly(x)) \supset Canfly(x)$, which changes the traditional nonmonotonic interpretation of "typical birds fly" to "birds that are known not to fly are the only birds that can't fly". Interpreted in this fashion sentences become statements of belief rather than typicality.

The difference between say default logic and autoepistemic logic is that default logic is tentative and thus defeasible, while autoepistemic logic contains context sensitive states whose evaluation may change as beliefs change. What follows from not knowing a proposition p will clearly change as soon as p has a truth value.

Moore defines a stable expansion E of a possible set of beliefs T an ideally rational agent can have based on a consistent set of sentences S as:

$$E = Th(S \cup \{L(P){:}P \in T\} \cup \{\neg L(P){:}P \notin T\}).$$

E is called a stable expansion of S. The consequences of this definition are that stable expansions contain all and only those formula supported by what is known to be true. Intuitively, this corresponds to the idea that different conclusions that are drawn with incomplete information are dependent upon what someone is prepared to believe. Moore's example is that, "I believe Richard Nixon is alive simply because if he was dead I surely would have heard about it."

Gelfond and Pryzmusinska's (1990a) reification for heterogenous multiple inheritance is based on autoepistemic logic. This is achieved by introducing the autoepistemic belief operator L to give a three-valued interpretation of sentences in the logical language (Moore, 1988). Their treatment considers the following translation between links in a network Γ and predicates:

$p \Rightarrow q \Leftrightarrow S_+(p,q)$	positive strict inheritance
$p \rightarrow q \Leftrightarrow D_+(p,q)$	positive default inheritance
$p \not\Rightarrow q \Leftrightarrow S_-(p,q)$	negative strict inheritance
$p \not\rightarrow q \Leftrightarrow D_-(p,q)$	negative default inheritance

Path entailment is defined by the following axiomatisation:

$holds(a,p)$ iff $a \rightarrow p \in \Gamma$
$L\ S_+(p,g) \wedge L\ holds(a,p) \supset hold(a,q)$
$L\ D_+(p,q) \wedge L\ holds(a,p) \wedge \neg L\ ab(a,p,q) \supset holds(a,q)$
$L\ D_-(p,q) \wedge L\ holds(a,p) \wedge L\ ab(a,p,q) \supset \neg holds(a,q)$

Additional axioms are required for the preference criteria for competing paths, namely specificity, inconsistency and preemption. Although this approach appears promising, it does have its drawbacks. Ginsberg (1990) points out that although authoepistemic logic provides a direct and local translation from multiple inheritance networks, the resulting axiomatisation of the global inheritance characteristics produces so large and complex an axiomatisation, even for reasonably trivial examples, that it is more reminiscent of a procedural definition than a truly declarative approach.

Having said this, Ginsberg's approach is itself not radically dissimilar in the sense that extensions for ambiguous multiple inheritance are presented in terms of abnormality rules which justify particular extensions. He defines specificity in terms of the causal relationship between normality assumptions and the abnormality of subclasses, and although considerably appealing, there is no guarantee that the axiomatisation will result in the appropriate intuitions for preemption, or produce the correct intuitive consequences under the addition of redundant links. It is, however, a purely local translation with natural primitives independent of most of the inheritance terminology.

5.6 Circumscriptive theories of inheritance

McCarthy and Hayes' (1981) initial treatment of inheritance as an application of circumscription attracted the criticism that it remained a user activity to choose the abnormality predicate to be minimised. McCarthy uses an object $(a,b\ldots)$, class $(p,q,\ldots)$, property $(x,y\ldots)$ description of inheritance (in a similar fashion as Gelfond and Pryzmusinska in section 5.5) to reify the representation of multiple inheritance.

$p \Rightarrow q$	$\equiv p \leq q$	specificity
$a \Rightarrow q$	$\equiv in(a,p)$	instantiation
$p \rightarrow x$	$\equiv ordinary(p,x)$	default property inheritance
$p \not\rightarrow x$	$\equiv ab(aspect(p,q,x))$	abnormal p w.r.t.
$p \Rightarrow q \rightarrow x$		inheriting x via q

Using this representation, default inheritance is subsequently axiomatised by

A1: $[ordinary(q,x) \wedge p \leq q \wedge \neg ab(aspect1(p,q,x))] \supset ordinary(p,x)$
A2: $[p \leq q \wedge q \leq r \wedge ordinary(q,not(x))] \supset ab(aspect1(p,q,x))$
A3: $[p \leq q \wedge q \leq r] \supset p \leq r$

Haugh (1988) points out that $A3$ presents a major problem in McCarthy's axiomatisation in that no provision is made for defeasible class membership relations. This is due to the object/class/property type-epistemology which distinguishes default object types by externalising them from the formalisation. Haugh proposes that $A3$ be replaced by the following axioms to allow for exceptions to class membership:

A3a: $[p \leq q \wedge q \leq r \wedge \neg ab(aspect1(p,q,r))] \supset p \leq r$
A3b: $[p \leq q \wedge q \leq r \wedge q \leq not(s)] \supset ab(aspect1(p,r,s))$

Haugh goes on to point out that McCarthy's initial theory will produce unintended models, and that a simple minimization of abnormalities will not produce the correct intuitive results. His reformulation defines the abnormality predicate *ab* in terms of the rules that generate abnormality. While McCarthy uses the predicates *aspect1* and *aspect2* to distinguish between class-to-class and object-to-class abnormality, Haugh's definition makes explicit the distinction within the *ab* predicate:

A1: $isa(X,Q) \equiv isa_x(X,Q) \vee \exists P\, isa(X,P) \wedge isa_x(P,Q) \wedge \neg ab(X,P,Q)$
A2: $ab(A,P,Q) \equiv ab_d(X,P,Q) \wedge ab_i(X,P,Q) \vee ab_c(X,P,Q) \vee ab_x(X,P,Q)$
A3: $ab_d(X,P,Q) \equiv isa(P,X) \wedge isa_x(P,Q) \wedge isa_x(X,not(Q))$
A4: $ab_i(X,P,Q) \equiv \exists R\, isa(X,R) \wedge ab_d(R,P,Q) \vee ab_c(R,P,Q) \vee ab_x(R,P,Q)$
A5: $ab_c(X,P,Q) \equiv \exists R\, isa(X,P) \wedge isa_x(P,Q) \wedge isa(X,R) \wedge isa_x(R,not(Q)) \wedge \neg ab_{dix}(X,R,not(Q))$
A5: $ab_{dix}(X,P,Q) \equiv ab_d(X,P,Q) \vee ab_x(X,P,Q) \vee \exists S\, isa(X,S) \wedge ab_d(S,P,Q) \vee ab_x(S,P,Q)$

Abnormality is thus defined in such a way that repeated computation for every query is avoided. Depending upon the desired sceptical/credulous interpretation (axioms *A4* & *A5*), abnormality can be propagated through the network. Abnormality is thus a "globally" defined concept which can be triggered by any number of local abnormality conditions. By grouping together the various abnormality conditions into a single abnormality predicate, the minimisation criteria is simplified and the many unnecessary abnormalities in McCarthy's proposal avoided. Minimisation of global abnormality corresponds to the results that would be achieved by local abnormality minimisation. When given a sceptical interpretation, Haugh's axioms produce a provable and unique model.

Further work in this vein by Krishnaprasnad (1989) and Krishnaprasnad and Kifer (1989) translates a network into a set of first-order sentences augmented with meta-level minimisation constraints based on prioritorised circumscription. Evidential support for individual default rules is incorporated into the object language and provides meta-knowledge about abnormality predicates. These can then be used to prioritorise the minimisation of the abnormality predicates in the circumscriptive theory. The allocation of evidential support values to defaults is done by a global examination of relative specificities of defeasible links. The general idea is that the evidential strength of defaults associated with subclasses are more significant than those of superclasses. The net effect is a Horn clause language where the implicit evidential order of the network is incorporated into the rule head. This embodies meta-knowledge about the priorities associated with abnormality predicates in the circumscriptive theory.

5.7 Preferential model treatments

Preferential model inheritance systems differ from default systems in that they have a radically different interpretation of default rules. In default logic, the conclusions drawn by the system can only be observed by examining the syntactic structure of a proof, whereas the preferential model approach has a natural and intuitive semantics but, in general, no proof theory. The general idea is that default rules are used to constrain the set of preferred, or most likely models of a situation. Each of these models is itself complete, while in default logic an extension only partially characterises a situation. Circumscription is a special case of preferential model theory where models are compared via the extension of selected predicates and a formula is preferentially entailed if it is true in ***all*** preferred models of the premise set.

Poole's (1985) work on inheritance is motivated by the problems of default logic treatments of multiple inheritance. In his approach, default knowledge is explicitly differentiated from first-order facts in the object language. Poole defines an independent semantic criteria by which default rules can be preferred over others on the basis of specificity. In doing so, the problems experienced by interacting default rules in Reiter's default logic are eliminated by the logic's semantics. This is done by selecting from the total set of defaults (Δ) a subset (D) which when combined with the first order sentences in the language (Γ) produce an extension (g), $\Gamma \cup D \vdash g$. The subset D of defaults is likened to a scientific theory, which explains the entailed conclusion g. In this sense the selection of the defaults to be included in D is an extra-logical activity beyond the semantics of the logic itself. The major distinction between the two approaches, apart from the explicit treatment of defaults in the object language[27] rather than Reiter's treatment of defaults as inference rules external from the object language, is the model theory. In default logic, extensions are maximally consistent sets of formulas which include the deductive closure of a base theory Γ and maximal instances of default inference rules. In Poole's approach, theories are minimal, only a minimal subset of defaults rules are chosen, which result in a consistent model, and the specificity criteria is incorporated into the logic's semantics rather than as a syntactic definition, *al la* Touretzky's (1986) inferential distance.

Loui's (1986) logic is similar to Poole's although some of the intuitions which result from reasoning explicitly with the defaults vary. Loui introduces additional criteria apart from default rule specificity by which explicit default rules can be preferred; namely, the criteria of "evidential preference", i.e., models which infer conclusions from maximal sets of defaults are preferred, and "directness" where models with shortest reasoning paths, which include common intermediate lines of reasoning, are preferred. The essential idea behind both Poole and Loui's approaches is minimal preferred models, where one or more preference relations can be implicitly defined over sets of competing models by examining the priorities on explicit default rules.

[27]A model is defined $g = (D,\Gamma)$ where D is the set of default statements and T the set of first order formulas. In this sense the defaults and first-order sentences are conceptually separated in the object language.

Selman and Kautz (1988) define a series of default systems progressing from a weak default system and progressively restraining that system to improve its performance. To produce a tractable proof procedure, the system is limited to representing system defaults (D) in Horn clause, form and ensuring that, when they are drawn, the defaults represent an acyclic multiple inheritance network. Given that a model $M = (D,T)$ is maximal, i.e., deductively closed, where Γ is the set of first-order formulas, then the models in this system can be ordered, $M \leqslant^{+} M'$ iff there exists some d in the set of default statements s.t. $(D \cup d, \Gamma) \vdash M'$. In other words, a minimal model is one in which all defaults are deductively satisfied. Defaults have the form $d = \alpha \rightarrow_{d} q$ and are applicable at a model M iff $M \vdash \alpha$ and d is not "blocked". Blocking is defined in terms of specificity and $d = \alpha \rightarrow_{d} q$ is said to be blocked at M if $\exists d' \in D$ where $d' = (\beta \cup \alpha) \rightarrow_{d} \neg q$ and $M \vdash (\beta \cup \alpha)$. The goal of default inference is to find a maximal model given a set of facts and an ordering on models defined by the default rules. Selman and Kautz give a polynomial algorithm for this system, and demonstrate that acyclic multiple inheritance can be represented in this way.

Stein's (1990) approach to model preferential entailment in multiple inheritance is novel, since it combines notions of path- and model-based treatments. The idea differs from the proof-theoretic or syntactic approaches of Horty (1990b) and Touretzky (1986) by using the notion of a preferential model semantics. It differs from the model-based approaches in that it does not rely on a local translation to a specific defeasible logic, but instead uses a credulous path definition, essentially a path-based definition, as the basis for a model. Stein believes that local translations of inheritance are undesirable since the inherent topology of the network is reduced to the semantics of the target logic whose ambiguity resolution strategy may not be appropriate for multiple inheritance.

Stein's approach is to treat credulous extension as translatable to propositional calculus. This is acceptable, despite the results discussed in section 5.1, since every credulous extension in itself can be seen as an unambiguous inheritance hierarchy. Relative to the definition of a credulous path as a propositional sentence, each model is both sound and complete. Given this model-theoretic interpretation of a credulous extension, the preference criteria is subsequently defined.

The preference criteria is specificity based. A credulous path $\phi = a \rightarrow p_1 \rightarrow \ldots \rightarrow p_n \rightarrow q_1 \rightarrow \ldots \rightarrow q_n \rightarrow s$ is preferred over the path $\varphi = a \rightarrow p_1 \rightarrow \ldots \rightarrow p_n \rightarrow r_1 \rightarrow \ldots \rightarrow r_n \rightarrow r_n \not\rightarrow s$, given that ϕ and φ agree on all edges $p_1 \rightarrow \ldots \rightarrow p_n$ and there is a subpath in φ which disagrees with ϕ w.r.t. a, i.e., $\varphi \vdash a \not\rightarrow s$ and $\phi \vdash a \rightarrow s$. A credulous extension is minimal under this ordering if there are no other credulous extensions which are preferred to it. In this approach, not only is the model-theoretic interpretation of a credulous extension simple, but the preferential criteria can easily be modified to accommodate off-path preemption as the basis for model preference.

5.8 Logics for inheritance—Summary

Etherington and Reiter (1983) used multiple inheritance with exceptions as an application by which to test default logic. They related default logic to inheritance by introducing a special exception link, which has no independent semantics. The point of the exception link is to alter the semantics of the default link so that semi-normal default rules can be created from the network.

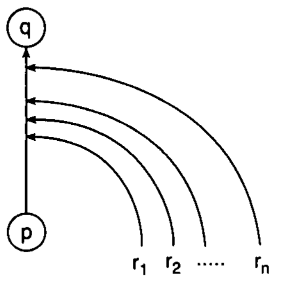

$= \Gamma_{24}$

In Γ_{24}, the exception links are introduced to explicitly allow a translation of inheritance links to semi-normal default rules that changes the semantics of bipolar multiple inheritance with exceptions by providing another level of partial order in the inheritance hierarchy, i.e. a partial order on competing default links is artificially introduced and there is absence of a modular translation from each link independent upon the context in which the link occurs (Touretzky et al., 1987b). Regardless of this criticism, the non semi-decidability of initial default logic treatments of inheritance made it an unlikely candidate for implementation. Its main contribution, arguably, was as a specification language.

Like Etherington and Reiter's (1983) early network translation, Padgham (1988, 1989a) introduces an inheritance lattice, essentially a binary node form, which provides a way of incorporating a partial order on default links via the modified lattice representation. The major difference between Etherington's and Padgham's approaches is that Padgham's adjustment to the modelling theory is made without a defeasible theory to substantiate the semantic changes to normative statements.

Having himself realised these difficulties with the prioritisation of default links, Etherington (1987b) was able to incorporate Touretzky's inferential distance ordering to overcome the problems, enabling them to define an order relation on their default expansions. Despite the subsequent elimination of the canellation link via an implicit order of defaults based on inferential distance (Etherington, 1987b), Etherington does not sufficiently examine the consequences to the proof procedure which such an implicit ordering would involve. Etherington (1988) went further to suggest that an inferential distance based inheritance reasoner could be used as a proof procedure for a restricted set of default logics, and this would thus overcome the intractability problem suffered by default logic in general. This suggestion was later refuted by Selman and Levesque (1989), who report that the computational complexity of any downward reasoner, regardless of design choices involving on- or off-path preemption or skeptical or credulous reasoning, is *NP*-hard.

The nonmonotonic properties of multiple inheritance with exceptions and the established relationship between multiple inheritance and default logic ensured the entry of other "nonmonotonists" into the debate. The nonmonotonists saw multiple inheritance with exceptions as an ideal vehicle for testing defeasible logics[28]. Etherington's claim that inheritance reasoners could be used to surmount intractability in proof procedures for nonmonotonic logics, once a suitable translation was found, may have also been an added incentive. Despite two decades of practical experimentation and application of inheritance hierarchies, surprisingly little progress has been made to establish a definitive model theoretics semantics. It seems clear that apart from the major question of the semantic validity of sceptical reasoning, most of the intuitive semantics of multiple inheritance have been discovered.

6 Implementation and tractability

One of the major motivations for path-based or algorithmic approaches to the inheritance problem are the poor performance characteristics of defeasible logics. Aside from the issue of representational adequacy, both default logic and circumscriptive inheritance treatments are undecidable in the first order case and intractable in the propositional case (Selman & Kautz, 1988). This is also true of autoepistemic and conditional logics.

Selman and Levesque (1989) prove that on-path, downward, credulous inheritance is *NP*-hard, and in doing so, demonstrate that Touretzy's conditioning algorithm, via the addition of redundant edges, is also intractable. This even holds for unambiguous networks. Generalising their results showed that all forms of downward reasoner are in fact intractable, while upward reasoners, regardless of the other design criteria[29], are polynomial in complexity. These results are reproduced in Table 3.

[28]Referring to nonmonotonic reasoning systems in general, i.e. autoepistemic logic, circumscription, conditional and default logics.

[29]Sceptical versus Credulous, on- versus off-path preemption.

Table 3. Selman and Levesque's (1989) tractability results

	Sceptical reasoning		*credulous reasoning*	
	off-path	*on-path*	*off-path*	*on-path*
Downwards	*NP*-hard	*NP*-hard	*NP*-hard	*NP*-hard
Upwards	*P*	*P*	*P*	*P*

The greatest difficulty facing credulous approaches to inheritance reasoners is that the multiple inheritance representation allows $O(2^{n-2})$ paths to be represented in an n node network. Figure 2 illustrates this.

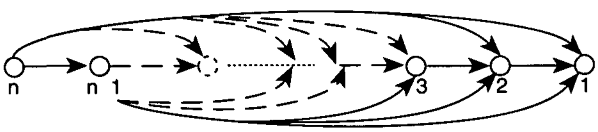

Figure 2 Complete networks = worst case multiple inheritance networks

Working with complete families of inheritance paths is thus intractable. On the other hand, complete networks, those which have direct connections between all node pairs in the network, are of little practical interest, and consequently worst case behaviour is seldom approached. The reason that worst case complexity results are misleading is that the network is generally sparse, that is to say, only a small number of the possible $n(n-1)/2$ network links are present in the network.

7 Conclusions

There is an essential tradeoff in multiple inheritance treatments between semantically sound but intractable approaches, i.e., those which use complete families of credulous extensions, and efficient semantic "approximations", namely sceptical reasoners. Given that the major semantic stumbling blocks of sceptical approaches to inheritance reasoning have been revealed, it is reasonable to assume that certain users will be prepared to forego semantic completeness in return for guaranteed tractability.

Another major issue is whether inheritance treatments should pursue a policy of locality, i.e., a local translation of networks into a formal representation via a direct definition of defeasible links[30], or alternatively, a path-based global analysis of the network to determine preferred entailments. The local translation approaches encourage incrementality and consider multiple inheritance as an application of more general classes of defeasible logics. The path-based treatments operate under the assumption that the topology of an inheritance network is crucial to its semantic interpretation, and that local translations cannot capture all the aspects of the intuitive inheritance semantics.

Treatments which rely on local translation, i.e., the model-based approaches, experience difficulties in capturing the semantics of path preemption, redundant link addition and sceptical reasoning, as well as suffering from complex axiomatisations. The path-based treatments cannot effectively produce local translations, a hindrance to incrementality, and suffer either from intractability or incompleteness. The fact that complete path-based treatments are intractable defeats one of the major arguments in their favour, that defeasible logics cannot tractably deal with the inheritance problem.

[30]See Krishnaprasad et al. (1989) and Ginsberg (1990) for a complete discussion of issues of locality.

It seems clear that there will be applications of inheritance where the semantically incomplete results obtained from sceptical reasoners will be unacceptable. This being so, the intractability of semantically correct multiple inheritance reasoning means that we increasingly have to discover syntactic network restrictions which will permit efficient computation. Perhaps on the basis of a computable defeasible theory and sufficiently strong syntactic restrictions, consensus inheritance intuitions can be grounded and a universal representation strategy will subsequently result.

Acknowledgements

This work has been supported by the Swedish Board of Technical Development (STU). Thanks to Patrik Doherty for his careful scrutiny and criticism offered at the time this survey was first written. Thanks also to an anonymous referee who helped restructure the survey.

Table 4. Author by author summary

Author	*Intuitive Semantics*	*Formal Semantics*	*Solution*
Bacchus	most p's are q's	$[q(x)\|p(x)]_x > b$	Probabilistic Lp
Boutiler	p's may may not be q's	$\neg(p \supset p) \wedge \neg(p \supset \neg q)$	conditional logic & minimal models
Doherty	p's typically have the property q	$(Lp \wedge M\,p_t \supset Dp_t) \wedge$ $(Lp_t \supset Lq)$	multi-valued logic & minimal models
Etherington	normally p's are q's	$\frac{p(x):q(x)}{q(x)}$	Default Logic
Geffner	infer q when p	if $p \rightarrow q$ then $p \vdash q$	probabilistic semantics
Gelfond	p's typically have the property q	*[31]	deductive axiomatisation in AEL
Ginsberg	p causes q	$p > q$ iff $p \in E$ $T \cup E \vdash q$, $T \cup E - \{p\} \nvdash q$	declarative semantics & preferred models
Grégoire	p is naturally a q	*	translation algorithm to stratified logic programs
Horty	p is naturally a q	*	inductive-definition of priorities & sequent calculus
Krishna- prasad	p is a q with confidence	priority on defaults w.r.t. others	multi-valued logic & prioritorised circumscription
McCarthy & Haugh	ordinarily p's have the property q	*	circumscription
Neufield	p is a partial set of q	$1 > P(q\|p) > P(q) > 0$	probabilistic
Padgham	p's are usually q's	model typicality by dualizing classes	lattice-based deductive definition
Pearl	p's typically have the property q	$P[q(x)\|p(x)] =$ HIGH	Probability Theory (ε-semantics)
Sandewall	any p is usually a q	*	deductive axiomatisation
Stein	arbitrary whether p is a q or not	*	—
Touretzky	p is naturally a q	*	Path-based proof-theoretic axiomatisation

[31]* Denotes no fixed interpretation on the normative statement $p \rightarrow q$ other than that p is more specific than q.

References

Bacchus, F, 1988. *Representing and reasoning with probabilistic knowledge*. PhD theseis, Department of Computer Science, University of Waterloo, Ontario, Canada.

Bacchus, F, 1989. "A modest, but semantically well founded, inheritance reasoner". In: *International Joint Conf. on Artificial Intelligence*, pp 1104–1109.

Belnap, N, 1976. "How a computer should think". In: G Ryle (ed.), *Contemporary Aspects of Philosophy*, pp 30–55. Oriel Press.

Belnap, N, 1977. "A useful four valued logic". In: J Dunn and G Epstein (eds.), *Modern Uses of Multiple Valued Logic*, pp 8–37. Reidel.

Bläsius, K, Hedück, U and Rollinger, C (eds.), 1990. *Sorts and Types in Artificial Intelligence: Lectures Notes on AI Vol 418*. Springer-Verlag.

Boutilier, C, 1989. "A semantical approach to stable inheritance reasoning". In: *International Joint Conf. on Artificial Intelligence*, 1134–1139.

Brachman, R, 1985. "I lied about the trees or defaults and definitions in knowledge representation". *AI Magazine*, Fall 80–93.

Buchanan, B and Shortliffe, E, 1984. *Rule-bases Expert Systems: The MYCIN Experiments*. Reading, MA: Addison-Wesley.

Borgida, A and Williamson, K, 1985. "Accommodating exceptions in databases and refining the schema by learning from them". In: *Proceedings of VLDB*. Stockholm.

Cross, C and Thomason, R, 1987. "Update and conditionals". In: *International Symposium on Methodologies for Intelligence Systems*.

Doherty, P, 1989. "A correspondence between inheritance hierarchies and a logics of preferential entailment". In: *International Symposium on Methodologies for Intelligent Systems*.

Doherty, P, 1991. *NML: A nonmonotonic formalism with explicit defaults*. Phd, Institutionen för Datavetenskap, University of Linköping.

Eklund, P, 1991. "Negotiating inheritance taxonomies in conceptual structures". In: *3rd Scandinavian Conference on AI*, pp 211–221. Roskilde, Denmark.

Etherington, D and Reiter, R, 1983. "On inheritance hierarchies with exceptions". In: *National Conference on Artificial Intelligence*, pp 104–108.

Etherington, D, 1987a. "Formalizing nonmonotonic reasoning systems". *Artificial Intelligence* **31** 41–85.

Etherington, D, 1987b. "More on inheritance hierarchies with exceptions". In: *National Conference on Artificial Intelligence*, pp 352–357.

Etherington, D, 1988. *Reasoning with Incomplete Information: Research Notes in AI*. See Mateo, CA: Morgan-Kaufmann.

Fahlman, S, 1979. *NETL: A System for Representing and Using Real-World Knowledge*. Cambridge, MA: MIT Press.

Lenzerini, M, Donini, F and Nardi, D, 1991. "Tractable concept languages". In: *International Joint Conf. on Artificial Intelligence*, pp 458–463.

Froidevaux, C and Kayser, D, 1988. "Inheritance in semantic networks and default logic". In: P Smets, E Mamdani, D Dubious and H Prade (eds.), *Non-Standard Logics for Automated Reasoning*, pp 179–212. New York, NY: Academic Press.

Gabbay, D, 1985. "Theoretical foundations for non-monotonic reasoning in expert systems". In: R Krzysztof (ed.), *Logics and Models of Concurrent Systems*, pp 439–457. Berlin: Springer-Verlag.

Geffner, H, 1989. *Default reasoning: Causal and conditional theories*. PhD thesis, Cognitive Systems Laboratory, Department of Computer Science, UCLA, CA.

Ginsberg, M, 1990. "A local formalization of inheritance: Preliminary report". In: *3rd International Workshop on Nonmonotonic Reasoning*, pp 119–1292. Lake Tahoe, CA.

Gelfond, M and Przymusinska, H, 1990. "Formalization of inheritance reasoning in autoepistemic logic". In: *3rd International Workshop on Nonmonotonic Reasoning*, 73–107. Lake Tahoe, CA.

Gelfond, M and Przymusinska, H, 1990. "Formalization of inheritance reasoning in autoepistemic logic". *Fundamenta Informaticae* **13** (4) 403–444.

Geffner, H and Verma, T, 1989. "Inheritance = chaining + defeat". In: *International Symposium on Methodologies for Intelligent Systems*.

Hägglund, S, 1989. "Iterative design and adaptive maintenance of knowledge-based office systems". *Office Information Systems: The Design Process*. Amsterdam: North-Holland.

Haugh, B, 1988. "Tractable theories of multiple defeasible inheritance in ordinary nonmonotonic logics". In: *National Conference on Artificial Intelligence*, pp 421–426.

Hayes, P, 1977. "In defence of logic". In: *International Joint Conf. on Artificial Intelligence*, pp 559–564.

Horty, J, 1990a. "A credulous theory of mixed inheritance". In: M Lenzerini (ed.), *Inheritance Hierarchies in Knowledge Representation*. Chichester: Wiley.

Horty, J, 1990b. "A skeptical theory of mixed inheritance". In: J Dunn and A Cupta (eds.), *Truth or Consequences*, pp 267–281. Amsterdam: Kluwer.

Horty, J and Thomasen, R, 1988. "Mixing strict and defeasible inheritance". In: *National Conference on Artificial Intelligence*, pp 427–432.

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

Krishnaprasad, T and Kifer, M, 1989. "An evidence-based framework for a theory of inheritance". In: *International Joint Conf. on Artificial Intelligence*, pp 1093–1098.

Krishnaprasad, T, Kifer, M and Warren, D, 1989. "On the declarative semantics of inheritance network". In: *International Joint Conf. on Artificial Intelligence*, pp 1099–1103.

Krishnaprasad, T, 1989. *The semantics of inheritance networks*. Phd thesis, Department of Computer Science, State University of New York, Stony Brook, NY.

Loui, R, 1986. "Defeat among arguments: A system of defeasible inference". Technical Report TR190, Department of Computer Science, University of Rochester.

McCarthy, J and Hayes, P, 1981. "Some philosophical problems from the standpoint of AI". In: B L Webber and N Nilsson (eds.), *Readings in AI*, pp 431–451. Los Angeles, CA: Morgan Kaufmann.

Moore, R, 1988. "Autoepistemic logic". In: P Smets, E Mamdani, D Dubious and H Prade (eds.), *Non-Standard Logics for Automated Reasoning*, pp 105–127. New York, NY: Academic Press.

Makinson, D and Schlechta, K, 1989. "Floating conclusions and zombie paths". In: *2nd International Workshop on Nonmonotonic Reasoning*, Sankt Augustin, FRG.

Makinson, D and Schlechta, K. 1991. "Floating conclusions and zombie paths: Two deep difficulties in the directly skeptical approach to defeasible inheritance nets". **48** 199–209.

Nebel, B, 1990. "Terminological reasoning is inherently intractable". *Artificial Intelligence* **43** (2) 235–249.

Padgham, I, 1988. "A model and representation for type information and its use in reasoning with defaults". In: *National Conference on Artificial Intelligence*, pp 409–414.

Padgham, L, 1989a. "Negative reasoning using inheritance". In: *International Joint Conf. on Artificial Intelligence*, pp 1086–1092.

Padgham, L, 1989b. *Non-monotonic inheritance for an object-orientated knowledge-base*. Phd thesis. Institutionen för datavetenskap, University of Linköping.

Pearl, J, 1987. "Deciding consistency in inheritance networks". Technical Report CSD 870053: R96, UCLA.

Pearl, J, 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan-Kaufman.

Poole, D, 1985. "On the comparison of theories: Preferring the most specific explanation". In: *International Joint Conf. on Artificial Intelligence*, pp 144–150.

Patel-Scheider, P, 1989. "Undecidability of subsumption in nikl*". *Artificial Intelligence* **39** 263–271.

Reiter, R and Crisuolo, G, 1981. "On interacting defaults". In: *International Joint Conf. on Artificial Intelligence*, pp 270–276.

Reiter, R, 1980. "A logic for default reasoning". *AI Journal* **13** 81–132.

Sandewall, E, 1986. "Nonmonotonic inference rules for multiple inheritance with exception". In: *IEEE*, pp 1345–1353.

Schubert, L, 1976. "Extending the expressive power of semantic networks". *Artificial Intelligence* **7** 163–198.

Schlechta, K, 1989. "Directly sceptical inheritance cannot capture the intersection of extension". In: *Workshop on Nonmonotonic Reasoning*, Sankt Augustin, FRG.

Shoham, Y, 1988. *Reasoning about Change*. Cambridge, MA: MIT Press.

Selman, B and Kautz, H, 1988. "The complexity of model-preference default theories". In: M Reinfrank, J de Kleer, M Ginsberg and E Sandewall (eds.), *2nd International Workshop on Nonmonotonic Reasoning*, pp 115–130. Springer-Verlag. Grassau, FRG.

Lehmann, D, Kraus, S and Magidor, M, 1990. "Nonmonotonic reasoning, preferential models and cumulative logics". **44** 167–207.

Selman, B and Levesque, H, 1989. "The tractibility of path-based inheritance". In: *International Joint Conf. on Artificial Intelligence*, pp 1140–1145.

Stein, LA, 1989. "Skeptical inheritance: Computing the intersection of credulous extensions". In: *International Joint Conf. on Artificial Intelligence*, pp 1153–1158.

Stein, I, 1990. "A preference-based approach to inheritance". In: *3rd International Workshop on Nonmonotonic Reasoning*, pp 233–245. Lake Tahoe, CA.

Thomason, R, Horty, J and Touretzky, D, 1987. "A calculus for inheritance in monotonic semantic nets". In: *International Symposium on Methodologies for Intelligent Systems*, pp 282–287. Elsevier.

Touretzky, D, Horty, J and Thomason, R, 1987. "A clash of intuitions: The current state of nonmonotonic multiple inheritance systems". In: *International Joint Conf. on Artificial Intelligence*, pp 476–482.

Touretzky, D, Horty, J and Thomason, R, 1987. "A skeptical theory of inheritance in nonmonotonic semantic networks". In: *National Conference on Artificial Intelligence*, pp 358–363.

Touretzky, D, 1984a. "Implicit ordering of defaults in inheritance systems". In: *International Joint Conf. on Artificial Intelligence*, pp 322–325.

Touretzky, D, 1984b. "The mathematics of inheritance systems". Technical Report CMU-CS-84-136, CMU.

Touretzky, D, 1986. *The Mathematics of Inheritance Systems: Research Notes in AI*. San Mateo, CA: Morgan-Kaufmann.

Touretzky, D and Thomason, R, 1988. "Nonmonotonic inheritance and generic reflexives". In: *National Conference on Artificial Intelligence*, pp 433–438.

Thomason, R and Touretzky, D, 1990. "Inheritance theory and networks with roles". Technical Report CMU-CS-90-139, CMU.