Published in *Proceedings of the Eighth International Symposium on Methodologies for Intelligent systems (ISMIS'93)*, Trontheim, Norway, June 1993, 132-141.

Combining Classification and Nonmonotonic Inheritance Reasoning: A First Step*

Lin Padgham
Department of Computer and
Information Science
Linköping University
S-581 83 Linköping, Sweden

Bernhard Nebel
German Research Center for
Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3
D-6600 Saarbrücken 11, Germany

Abstract

The formal analysis of semantic networks and frame systems led to the development of *nonmonotonic inheritance networks* and *terminological logics*. While nonmonotonic inheritance networks formalize the notion of default inheritance of typical properties, terminological logics formalize the notion of defining concepts and reasoning about definitions. Although it seems to be desirable to (re-)unify the two

*This work has been supported by the the Swedish National Board for Technical Development (STU) under grant # 9001669, by the Swedish Research Council for Engineering Sciences under grant # 900020, by the German Ministry for Research and Technology (BMFT) under research contract ITW 8901 8, and by the European Community as part of the ESPRIT Working Group DRUMS-II.

approaches, such an attempt has not been made until now. In this paper, we will make a first step into this direction by specifying a nonmonotonic extension of a simple terminological logic.

1 Introduction

The formal analysis of early semantic network and frame formalisms led to the development of two different families of knowledge representation formalisms, namely, *nonmonotonic inheritance networks* [5] and *terminological logics* [12]. Nonmonotonic inheritance networks formalize the idea of default inheritance of typical properties. Terminological logics aim at formalizing the idea of defining concepts and reasoning with such definitions, for instance, determining *subsumption* relationships between concepts and *instance* relationships between objects and concepts—two kinds of inferences we will collectively refer to as *classification*.

Although these two forms of representation and reasoning may seem to be incompatible [2], it would of course be desirable to combine them. From the point of view of nonmonotonic inheritance networks, it would be interesting to have a richer description language for specifying classes and properties and to add the ability of *classifying* objects as belonging to some class. From the point of view of terminological logics, it is desirable to add forms of reasoning that deal with uncertain information. In fact, Doyle and Patil [3] argue that a representation system without such a facility is useless.

There are proposals to integrate some form of default inheritance in terminological logics since 1981 (see [14, 12]) and some terminological representation systems support forms of nonmonotonic inheritance, which appear to combine the two modes of reasoning in a "naive" way (e.g. [12]), however, leading to problems similar to the infamous "shortest path inference," as we will see in Section 3.

An attempt to combine classificatory reasoning and nonmonotonic inheritance that avoids the latter problem has been made by Horty and Thomason [4]. Although this approach comes closest to our intention of combining nonmonotonic inheritance and classification, there are some problems, for instance, the "zombie path" problem [7], the lack of an algorithm, and the computational intractability of the approach.

More recent approaches combine classificatory and nonmonotonic reasoning by integrating default logic into terminological logics [1]—without using specificity for conflict resolution, though—or they employ a form of preference semantics [13].

We will base our combination of inheritance reasoning and classification on the nonmonotonic inheritance reasoning approach by Padgham [10], which avoids the above mentioned shortcomings. In the following sections we introduce a restricted terminological logic extended by defaults, and discuss how the inheritance theory in [10] can be extended to include classification.

2 A Common Representational Base

In order to describe our approach, we first introduce a representation formalism that can be conceived as a restricted *terminological logic*.

We start with a set $\mathbf{A}$ of *atomic concepts* (denoted by $A, A', \ldots$) and a set $\mathbf{F}$ of *features* (denoted by $F, F', \ldots$) that are intended to denote single-valued roles. Additionally, we assume a set $\mathbf{V}$ of *values* (denoted by v, v') that are intended to denote atomic values from some domain. Based on this, *complex concept expressions* (denoted by C, C') can be built:

$$C \;\rightarrow\; \top \mid \bot \mid A \mid C \sqcap C' \mid F\colon v.$$

In order to *define* new concepts completely or partially, *terminological axioms* (denoted by θ) are used. *Assertions* (denoted by α) are employed to specify properties of objects ($x, y, z, \ldots \in \mathbf{O}$):

$$\theta \;\rightarrow\; A \sqsubseteq C \mid A \doteq C,\; \alpha \;\rightarrow\; x\colon C \mid F(x) \doteq v.$$

Knowledge bases are sets of such terminological axioms and assertions.

The semantics of this language is given in the usual set-theoretic way. An interpretation $\mathcal{I}$ is a tuple $\langle \mathcal{D}, \mathcal{V}, \cdot^{\mathcal{I}} \rangle$, where $\mathcal{D}$ and $\mathcal{V}$ are arbitrary non-empty sets that are disjoint, and $\cdot^{\mathcal{I}}$ is a function such that

$$\cdot^{\mathcal{I}}\colon (\mathbf{A} \rightarrow 2^{\mathcal{D}}) \cup (\mathbf{F} \rightarrow 2^{(\mathcal{D} \times \mathcal{V})}) \cup (\mathbf{V} \rightarrow \mathcal{V}) \cup (\mathbf{O} \rightarrow \mathcal{D})$$

where we assume that the relation denoted by a feature is a *partial function* and that values and object identifiers satisfy the *unique name assumption*. Interpretations are extended to complex concept expressions in the usual way, e.g., $(C \sqcap C')^{\mathcal{I}} = C^{\mathcal{I}} \cap C'^{\mathcal{I}}$ and $(F\colon v)^{\mathcal{I}} = \{d \in \mathcal{D} \mid (d, v^{\mathcal{I}}) \in F^{\mathcal{I}}\}$.

An interpretation is called a *model* of a knowledge base, if all terminological axioms and assertions are *satisfied* by the interpretation in the obvious way, e.g., $A^{\mathcal{I}} = C^{\mathcal{I}}$ for $\theta = (A \doteq C)$. The specialization relationship between concepts (also called *subsumption*) and the *instance relationship* between object identifiers and concepts are defined in the obvious way. A concept C is subsumed by C', written $C \preceq C'$ iff $C^{\mathcal{I}} \subseteq C'^{\mathcal{I}}$ for all models $\mathcal{I}$ of the knowledge base. An object x is an instance of a concept C, written $x\colon C$, iff for all models $\mathcal{I}$ it holds that $x^{\mathcal{I}} \in C^{\mathcal{I}}$.

In order to express that an instance of a concept C typically has some additional properties, the syntax of terminological axioms is extended as follows:

$$\theta \rightarrow A \sqsubseteq C/D_1, \ldots, D_n \mid A \doteq C/D_1, \ldots, D_n,$$

where the D_i's are again concept expressions. These "default properties" do not influence the set-theoretic interpretation of concepts, but are intended to denote that an instance of A typically has the additional properties D_i. In terms of Padgham's [10] type model, given an axiom $A \sqsubseteq C/D_1, \ldots, D_n$, C represents the *core* of a type, while $C \sqcap D_1 \sqcap \ldots \sqcap D_n$ is the *default* of a type.

Using our simple representation formalism, we could classify concepts in the TBox, compute instance relationships between objects and concepts, and separately apply default inheritance in order to derive typical information about objects. In fact, this loose combination of classification and default inheritance was used profitably in a medical diagnosis application [15].

The network language we will use contains strict inheritance links "$\Rightarrow$", strict negative links $\not\Rightarrow$", and default inheritance links "$\rightarrow$". In addition to the usual kind of nodes, depicted by a letter, we also allow for *defined* nodes, depicted by an encircled letter. The latter nodes are assumed to be defined by the conjunction of all nodes that are reachable by a single strict positive link. As an example let us consider the following small knowledge base (inspired by [2]):

$$
\begin{array}{rcl}
\texttt{Elephant} & \sqsubseteq & \texttt{Mammal/legs:}\,4, \texttt{color:}\,\texttt{grey} \\
\texttt{Hepatitis-Elephant} & \doteq & \texttt{Elephant} \sqcap \texttt{infected-by:}\,\texttt{Hepatitis/color:}\,\texttt{yellow} \\
\texttt{Yellow-Elephant} & \doteq & \texttt{Elephant} \sqcap \texttt{color:}\,\texttt{yellow} \\
x\colon\texttt{Elephant} & & \\
\texttt{infected-by}(x) & \doteq & \texttt{Hepatitis}
\end{array}
$$

Using the abbreviations M, E, H, and Y for `Mammal`, `Elephant`, `Hepatitis-Elephant`, and `Yellow-Elephant`, respectively, and g, h, l, y for `color`: `grey`, `infected-by`: `Hepatitis`, `legs`: 4, and `color`: `yellow`, respectively, the network diagram corresponding to our small knowledge base would look like as in Figure 1.

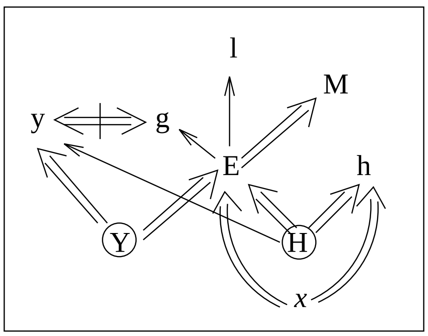

Figure 1: Shortest Path Problems

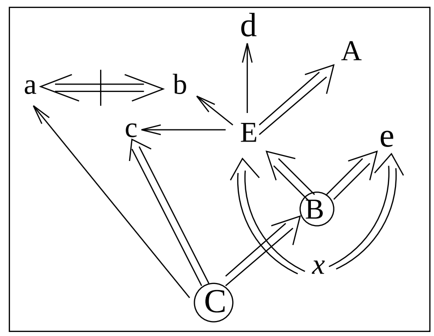

Figure 2: Example Network

3 Some Problems Combining Defaults with Classification

Using concept specificity for resolving conflicts among contradicting typical properties seems to be natural and desirable. Indeed, most proposals or already implemented systems seem to prefer this kind of conflict resolution. MacGregor, for instance, integrated a facility for "specificity-based defaults" [6, p. 393] into LOOM. The proposal by Pfahringer [12] is also an effort in this direction, employing a form of skeptical inheritance as defined in the area of nonmonotonic inheritance reasoning. From the limited descriptions of such approaches in the literature they appear to combine the two modes of reasoning in what we will call a "naive" way, which can be described as follows. Given an object x and a description D_0 of x, we first determine the set of most specialized concepts S such that x is an instance of all of the concepts in S. Based on this we determine additional typical properties of x using some inheritance strategy, which gives us a new (more specialized) description D_1, and we start the cycle again. We stop when a fixpoint is reached, i.e., D_i is equivalent to D_{i-1}.

The main problem with the "naive" approach is that it leads to results resembling the infamous shortest path inference. This can be illustrated by considering figure 2. If we begin by classifying x we get B. Default reasoning then gives $b, c, d. \neg a$ and a further round of classification gives C. we would now want by default to believe a, but this is blocked because we believe b and $\neg a$. However we observe that the default belief in a comes from a more specific type (C), than the default belief in b (which comes from E). We would therefore prefer to believe a than b. However we have previously commited to b because we reached it first.

4 A Default Inheritance Reasoning Framework

In this section we develop a formal framework for default inheritance reasoning. We will then generalize this in the following section so that it becomes a framework for combined classification and default reasoning. The framework that we develop is based on that presented in [10, 11]. The theory is very close to the skeptical inheritance theory of Horty *et al* [5] in terms of the conclusions reached[1], but instead of working with constructible, pre-empted and conflicted paths, we work with notions of *default assumptions*, *conflicting assumptions* and *modification of assumptions* in order to resolve conflicts.

Given some initial information and an inheritance net, we first assemble all the default assumptions that may be possible, given this start point. We then find all the pairwise conflicting assumptions, and resolve the conflicts—starting with most specific nodes—by modifying one or both of the assumptions. Finally we add all our modified (and now consistent) assumptions together to obtain our set of conclusions.

Our formalization is based on an inheritance network, Γ, which is derivable directly from terminological axioms and assertions as defined in Section 2, and *labellings* which are mappings from the nodes in the network (the set $\mathbf{N}_\Gamma$) to values in the set $\{0, 1, -1, k\}$. The intuitive interpretation of such a labelling L is an information state concerning a hypothetical object where $L(X) = 1$ means that the object is an instance of the concept X, $L(X) = -1$ means that the object is not an instance, $L(X) = 0$ means there

[1]It does not however have the "zombie path" behavior criticized by Makinson and Schlechta [7].

is no information concerning the instance relationship and $L(X) = k$ means there is contradictory information. The set $\{0, 1, -1, k\}$ forms a lattice w.r.t. information content such that $0 \leq 1 \leq k$, and $0 \leq -1 \leq k$. Similarly, the set of all labellings forms a lattice based on this ordering. In particular the *join* of two labellings, written $L_1 \sqcup L_2$, corresponds to the combination of the information content. The special labelling $\underline{0}$ is the labelling with all labels 0.

A labelling L is said to be *consistent* if it does not contain any node with a value of k. A pair of labellings is said to be *compatible* if their join is consistent, and *weakly compatible* if their join does not introduce any new inconsistency not present in one of the individual labellings.

Definition 1 *There is a* **strict positive path** *from X to Y in* Γ, *written* $X \Rightarrow \sigma \Rightarrow Y$, *iff exists W,Z:* $(X{=}W \vee [X \Rightarrow \sigma \Rightarrow W] \in \Gamma) \wedge [Z \not\Leftrightarrow W] \in \Gamma \wedge (Y{=}Z \vee [Y \Rightarrow \sigma \Rightarrow W] \in \Gamma)$

Definition 2 *There is a* **strict negative path** *from X to Y in* Γ, *written* $X \not\Rightarrow \sigma \Leftarrow Y$, *iff* $[X \not\Leftrightarrow Y] \in \Gamma \vee (\exists W\!: [X \not\Leftrightarrow W] \in \Gamma \wedge [Y \Rightarrow \sigma \Rightarrow W] \in \Gamma)$.

We define two particular kinds of labellings—*core labellings* (written X_c) and *default labellings* (written X_d) for a node X. A core labelling represents the necessary information for a node, while a default labelling for a node X represents the information typically associated with X.

Definition 3 *A core labelling for X (w.r.t.* Γ*), written* X_c*, is the minimal labelling which fulfills the following:*
$X_c(X) \geq 1$*; and for all Y*
$([X \Rightarrow \sigma \Rightarrow Y] \in \Gamma) \rightarrow (X_c(Y) \geq 1) \wedge ([X \not\Rightarrow \sigma \Leftarrow Y] \in \Gamma) \rightarrow (X_c(Y) \geq -1)$

Definition 4 *A* **default labelling** *for X (w.r.t.* Γ*), written* X_d*, is the minimal*[2] *labelling which fulfills the following:*
$X_d \geq X_c$*; and for all Y* $([X \rightarrow Y] \in \Gamma) \rightarrow (X_d \geq Y_c)$

Referring back to Figure 1, the core labelling for `Hepatitis-Elephant` would have values of $\{H = 1,\ h = 1,\ E = 1,\ M = 1 \textit{ all else} = 0\}$. Its default

[2]The ordering over labellings is the obvious one, given the ordering over node values.

labelling would contain $\{H = 1,\ h = 1,\ E = 1,\ M = 1,\ y = 1,\ g = -1,\ \mathit{all\ else} = 0\}$.

Referring again to Figure 1, we may wish to block that part of E_d (default elephant assumption) which concludes g (grey), but allow a modified assumption which concludes l (four legs). On the basis of default labellings we introduce the notion of *modified assumption* (written $X_{d'}$).

We define a *correct* modified assumption which intuitively allows removal only of entire branches from the full default assumption. Correctness ensures both consistency w.r.t. the network and also that (potentially) dependent properties are treated together.

Figure 3 gives some examples of correct and incorrect modified assumptions.[3]

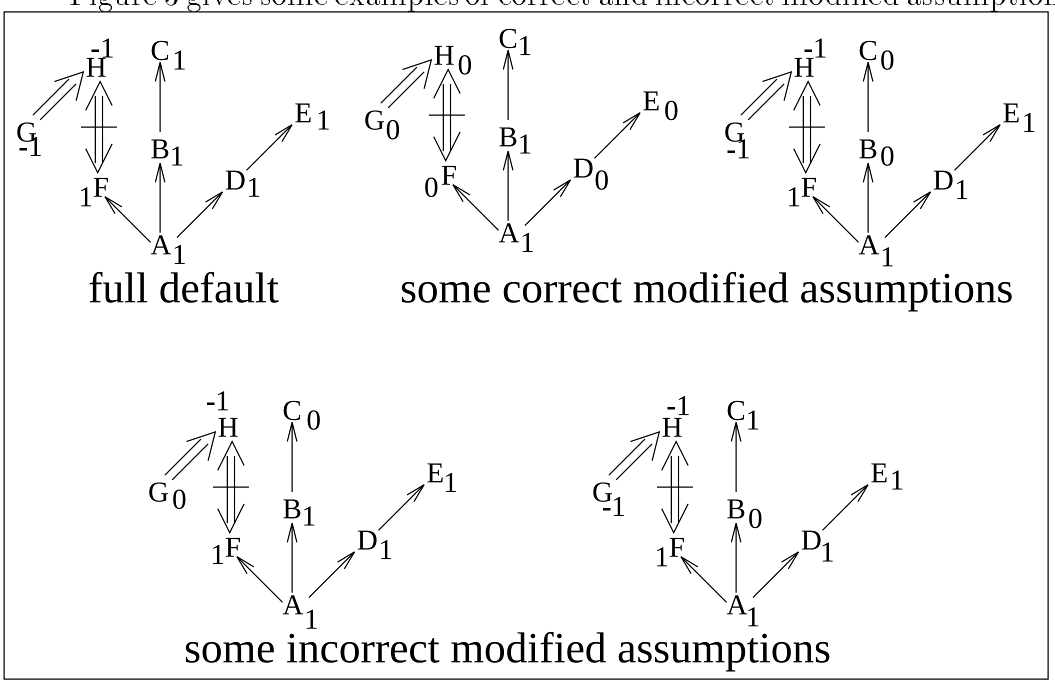

Figure 3: Correct and Incorrect Assumption Modifications

Definition 5 *A modified assumption, $X_{d'}$, is* **correct** *iff the following conditions hold for all Y in Γ:*

1. $([X \rightarrow Y] \in \Gamma) \rightarrow ((X_{d'}(Y) = 1) \rightarrow\ X_{d'} \geq Y_c$

[3]The concept of correctness is further motivated and explained in [10, p. 188–189], where it is dealt with as two separate concepts - groundedness and consistency.

2. $X_{d'}(Y) = 1 \rightarrow (\exists Z\colon [X \rightarrow Z] \in \Gamma \wedge Z_c(Y) = 1)\ \ X_c(Y) = 1$

3. $X_{d'}(Y) = -1 \rightarrow\ (\exists Z\colon [X \rightarrow Z] \in \Gamma \wedge X_{d'}(Z) = 1 \wedge Z_c(Y) = -1) \vee X_c(Y) = -1$.

4. $X_{d'} \geq X_c$.

When an assumption is modified it is always modified with respect to some other information with which it is in conflict. We thus introduce the notion of a modified assumption as a pair of labellings consisting of the default assumption labelling for the node and a *preference labelling* for the node (written P). The preference labelling captures all of the information which is to be preferred over the default assumption at that type. While it can in principal be an arbitrary labelling the preference labelling will for all interesting theories depend on both the type network and the initially given information. There is no constraint on the preference labelling to be consistent.

The value of the preference labelling for a node determines the modified assumption for that node. If the preference labelling for a node X is not weakly compatible with X_c (indicating preferred disbelief in the concept), then the modified assumption will be empty. Otherwise the modified assumption is a labelling between the core and the default, w.r.t. information content.

Definition 6 *A* **modified assumption** $X_{d'} = (X_d, P)$ *is* $\underline{0}$ *iff* X_c *is not weakly-compatible with* P*; otherwise* $X_{d'}$ *is the maximal labelling that is weakly-compatible with* P*, is a correct modification of* X_d *and* $X_c \leq X_{d'} \leq X_d$.

By joining a set of modified assumption labellings for a given network we can obtain a conclusion labelling for that network. We want to ensure that the preference labellings modify the default assumptions sufficiently to remove all conflicts so that we can obtain a consistent conclusion labelling. Preference labellings will be determined by the structure of the taxonomy together with the initial information, using principles such as specificity.

Each node in the network has its own preference labelling. We call this collection of preference labellings a *preference map*, written Θ. For a given network Γ, and a given initial labelling ψ, the preference map provides a preference labelling Θ_X for each node X in $\mathbf{N}_\Gamma$. Different inheritance theories can be compared with respect to the characteristics of their preference map. We will first characterize what we call a *well-formed* preference map, which can then be used as a base for defining a preference map for different kinds of theories, e.g. skeptical and credulous preference maps.

The characteristics that we capture in the definition of a well-formed preference map are that initial information is always preferred over default assumptions, more specific information is always preferred over less specific, unless the more specific information is unsupported, and that only supported (or reachable) modified assumptions are non-empty.

Definition 7 *A preference map* Θ *is* **well-formed** *for a network* Γ *and an initial labelling* ψ*, iff the following conditions are satisfied:*

1. $\Theta_X(X) < k$*, for all* $X \in \mathbf{N}_\Gamma$*,*
2. $\Theta_X \geq \psi$*, for all* $X \in \mathbf{N}_\Gamma$*,*
3. *there exists a strict partial ordering* $\ll$ *such that for all* $X \in \mathbf{N}_\Gamma$*: if* $\Theta_X(X) \leq 1$*, then* $\psi(X) = 1$ *or there is a* $Y \in \mathbf{N}_\Gamma$ *s.t.* $Y \ll X$ *and* $Y_{d'}(X) = 1$*,*
4. *if* X *is more specific*[4] *than* Y*, then* $X_{d'} \leq \Theta_Y$*.*

Θ^0 denotes the minimal well-formed preference map.[5]

To characterize a skeptical preference map we require in addition to well-formedness that each pair of modified assumptions are either compatible under well-formed preference, or that the preference labelling for each includes the other (forcing modification of each w.r.t. the other).

Definition 8 *A preference map* Θ *is* **skeptical** *for a network* Γ *and an initial labelling* ψ*, iff it is well-formed and for all* $X, Y \in \mathbf{N}_\Gamma$*:* $(X_d, \Theta^0_X) \sqcup (Y_d, \Theta^0_Y)$ *is consistent or* $(\Theta_X \geq Y_d$ *and* $\Theta_Y \geq X_d)$*.*

[4]For an exact definition of the notion of specificity used see [10, p. 142].

[5]Proof of the existence of a unique minimal well-formed preference map is due to Ralph Rönnquist, and can be found in [10] (where it is called a revision function).

5 Integrating Classification into the Framework

As we saw in section 3 we should not simply interleave correct classification and correct default reasoning as this will lead to a certain arbitrariness in the results, and will give problems analogous to "shortest path" problems found in early approaches to default inheritance reasoning. We therefore take the approach of defining a single theory which includes both classificatory reasoning and default reasoning.

Condition 1 of Def. 7 for a well-formed preference map simply states that a preference labelling should not be inconsistent regarding the preference of the node for which it is a preference labelling. This is not affected by classification.

Condition 2 of Def. 7 captures that initial information should be preferred over all default assumptions. This is also a criteria which is clearly applicable to combined classificatory/default reasoning.

Condition 4 of Def. 7 says that we prefer assumptions associated with more specific, rather than less specific assumptions and also seems appropriate to retain unchanged.

The final condition of well-formedness, (condition 3) has to do with ensuring that a default assumption is empty unless we independently from it (and its results) believe in the base concept. Looking at Figure 4, and starting with information F, we clearly would not want to make any default assumptions regarding, for example, C or X.

In the default inheritance reasoning the support in the ordering of condition 3 of well-formedness, is shown by a labelling of 1 on a node in some "earlier" modified assumption. However if we include classification as a valid means of reaching a conclusion then support may come from not only single assumption(s) but from a set of assumptions, which, taken together provide the "evidence" for believing that type. In order to capture this formally, we define the notion of *support*.

Definition 9 *A set of labellings* Π **supports** *a node* X *iff for some labelling* $L \in \Pi$: $L(X) = 1$ *or for all* $Z \in \mathbf{N}_\Gamma$: $((X_c(Z) = 1$ *implies* Π *supports* $Z)$ *and* $(X_c(Z) = -1$ *implies that for all* $L \in \Pi$: $L(Z) \leq -1))$.

Note that the support required for default reasoning is simply a special case of this definition of support. We can now rewrite the third condition of well-formedness as follows:

3. *there is a strict partial ordering* $\ll$ *on the nodes in* $\mathbf{N}_\Gamma$ *such that* $\forall X \in \mathbf{N}_\Gamma$*: if* $\Theta_X(X) \leq 1$*, then* $\{\psi\}$ *supports* X *or there exists* Π *s.t.* Π *supports* X *and for all* $L \in \Pi$ *there exists a node* $Y \in \mathbf{N}_\Gamma$ *s.t.* $L = Y_{d'}$ *and* $Y \ll X$.

The additional criterion for a skeptical preference map ensures that any ambiguous conflicts which remain following application of specificity for conflict resolution will result in bilateral modification of the conflicting assumptions. This appears to be equally applicable to combined classification/default reasoning as it is to pure default inheritance reasoning. To illustrate the principle captured here we observe Figure 5. In the left-hand network, the default assumptions at E and B are both modified to avoid concluding $G, \neg F$ or $F, \neg G$ respectively.

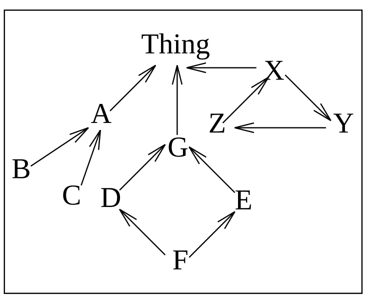

Figure 4: Support

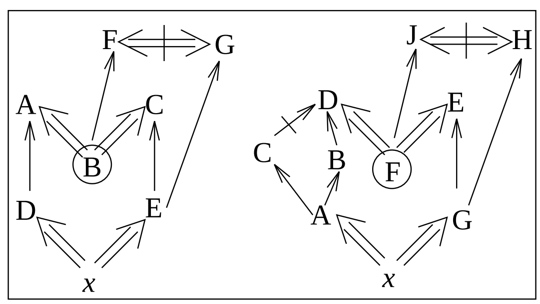

Figure 5: Ambiguous Conflicts

In the right hand network, the default assumptions at C and B are modified leading to no conclusion regarding D, and consequently no classification of F. However, because $\{B, C\}$ is an ambiguous conflict, the classification F would be made in some credulous extension, and we therefore allow it to cause modification to the assumption at E regarding H. Thus the conclusion for this figure will be A, B, C, G, E. We note that this is different than the extension given by Horty's method which will also include H and $\neg J$. H and $\neg J$ are not in the intersection of credulous extensions, and thus we would argue that they should not be in the skeptical extension. This difference is a result of the different treatment of ambiguous conflict in the two methods, where our approach is what has been referred to as "ambiguity propagating" in order to avoid the oddities of "zombie paths" [7].

6 Discussion

We have shown how a minor change to a theory for default inheritance gives a theory of combined classification and default inheritance for a very restricted terminological logic. Whilst this language may be too restricted for real applications it allows us to obtain a clearer understanding of the complex interaction between default inheritance and classification. This provides a firm basis on which we can begin to experiment with the addition of some greater expressivity. It may well be necessary to limit the expressivity of terminological languages with defaults, not so much because of tractability problems as most terminological languages are already intractable [8], but because of problems with "conceptual complexity". However there are certainly some applications which require defaults but whose other requirements on expressivity are limited [15]. The approach described in this paper provides a start for investigating languages and associated reasoning mechanisms for such applications.

The theory developed here is based on a skeptical inheritance theory with a ploynomial algorithm. Considering the minor nature of the change required to incorporate classification into this theory, the algorithm should also be directly modifiable to include classification.

References

[1] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. In Nebel et al. [9], pages 306–317.

[2] R. J. Brachman. 'I lied about the trees' or, defaults and definitions in knowledge representation. *The AI Magazine*, 6(3):80–93, 1985.

[3] J. Doyle and R. S. Patil. Two theses of knowledge representation: Language restrictions, taxonomic classification, and the utility of representation services. *Artificial Intelligence*, 48(3):261–298, Apr. 1991.

[4] J. F. Horty and R. H. Thomason. Boolean extsionsins to inheritance networks. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 633–639, Boston, MA, Aug. 1990. MIT Press.

[5] J. F. Horty, R. H. Thomason, and D. S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. In *Proceedings of the 6th National Conference of the American Association for Artificial Intelligence*, pages 358–363, Seattle, WA, July 1987.

[6] R. MacGregor. The evolving technology of classification-based knowledge representation systems. In J. F. Sowa, editor, *Principles of Semantic Networks*, pages 385–400. Morgan Kaufmann, San Mateo, CA, 1991.

[7] D. Makinson and K. Schlechta. Floating conclusions and zombie paths: Two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets. *Artificial Intelligence*, 48(2):199–211, Mar. 1991.

[8] B. Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43:235–249, 1990.

[9] B. Nebel, W. Swartout, and C. Rich, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference*, Cambridge, MA, Oct. 1992. Morgan Kaufmann.

[10] L. Padgham. *Non-Monotonic Inheritance for an Object-Oriented Knowledge Base*. PhD thesis, University of Linköping, Linköping, Sweden, 1989. Linköping Studies in Science and Technology. Dissertations No. 213.

[11] L. Padgham. Defeasible inheritance: A lattice based approach. *Computers and Mathematics with Applications*, 23(6-9):527–541, 1992. Special Issue on Semantic Nets.

[12] P. F. Patel-Schneider, B. Owsnicki-Klewe, A. Kobsa, N. Guarino, R. MacGregor, W. S. Mark, D. McGuinness, B. Nebel, A. Schmiedel, and J. Yen. Term subsumption languages in knowledge representation. *The AI Magazine*, 11(2):16–23, 1990.

[13] J. Quantz and V. Royer. A preference semantics for defaults in terminological logics. In Nebel et al. [9], pages 294–305.

[14] J. G. Schmolze and R. J. Brachman, editors. *Proceedings of the 1981 KL-ONE Workshop*, Cambridge, MA, 1982. Bolt, Beranek, and Newman Inc. BBN Report No. 4842.

[15] T. Zhang and L. Padgham. A diagnosis system using inheritance in an inheritance net. In *Proceedings of the Fifth International Symposium on Methodologies for Intelligent systems*, Knoxville, TN, Oct. 1990. North-Holland.