On Principles and Problems of Defeasible Inheritance

Karl Schlechta
DFKI, Stuhlsatzenhausweg 3,
D-6600 Saarbrücken, Germany
and
David Makinson
Les Etangs B2, La Ronce,
F-92410 Ville d'Avray, France

December 10, 1992

Abstract

We have two aims here: First, to discuss some basic principles underlying different approaches to Defeasible Inheritance; second, to examine problems of these approaches as they already appear in quite simple diagrams. We build upon, but go beyond, the discussion in the joint paper of Touretzky, Horty, and Thomason: A Clash of Intuitions, [THT].

1 INTRODUCTION

All diagrams are collected in the Appendix. We shall limit the discussion to purely defeasible finite acyclic inheritance nets. These are (finite acyclic) graphs, with two kinds of nodes and arrows (links). The nodes stand for objects or classes of objects. To simplify matters, we assume in our theoretical discussions to have only class-type nodes. The positive arrows $x \to y$ mean something like "normally, x's are y's", or "most x are y", the negative ones $x \not\to y$ "most x are not y". We consider a net Γ as information from which to draw conclusions of two kinds: First, we may say that Γ permits some "line of reasoning", i.e. a "path" of concatenated arrows of Γ, like σ: $u \to x \not\to y$ in diagram Γ_1, second, that Γ permits a result, like "most u are not y" here, if this is the conclusion of a permitted line of reasoning. Permitted paths and results will also be called "valid" or "accepted" in Γ, sometimes written $\Gamma \models \sigma$ etc. To differentiate from arrows, positive results will be noted xy, negative ones $x\overline{y}$. Thus, the conclusions of a net are very meagre in language: not even conjunction is permitted. The intuitive reading immediately shows that negative arrows should be permitted only at the end of an accepted path. Such paths with at most one negative arrow (at the end), will be called potential paths. In the absence of any conflicts, all potential paths are clearly valid. The problem is to single out the valid paths among the potential paths by suitable definitions, when there are conflicts.

We shall examine several definitions of acceptability of paths and their results for defeasible inheritance nets. In Part 2, we try to shed light on some of the fundamental differences between various definitions, in Part 3, we discuss problems specific to some approaches, in Part 4, we discuss some horizontal and vertical coherence properties. The principles common to all the definitions are: 1) No ex falso quodlibet, contradictions should be kept local; 2) More specific information should win in case of conflict; 3) Lacking differences in specificity ("unbiased conflict"), contradictions should be handled fairly. In addition, to get started, we shall assume that all arrows in Γ are accepted paths in Γ of length 1. These principles are still liberal enough to permit a multitude of approaches. For lack of space, some familiarity with defeasible inheritance has to be assumed - see e.g. [H], [HTT], [THT], [T].

2 FUNDAMENTAL DIFFERENCES

2.1 Extension-Based vs. Directly Skeptical Definitions

As this distinction has already received detailed discussion in the literature, we can be very brief here. An extension of a net is essentially a maximally consistent and in some appropriate sense reasonable subset of all its potential paths. This can of course be presented either as a liberal conception (focussing on individual extensions) or as a skeptical one (focussing on their intersection - or more accurately, as discussed in section 2.5 below, the intersection of their conclusion sets). The seminal presentation is that of [T], as refined by [Sa]. The directly skeptical approach seeks to obtain a notion of skeptically accepted path and conclusion, but without detouring through extensions. Its classic presentation is that of [HTT]. Even while still searching for fully adequate definitions of either kind, we may use the former approach as a useful "control" on the latter. For if we can find an intuitively possible and reasonable extension supporting a conclusion $x\overline{y}$, whilst a proposed definition for a directly skeptical notion of legitimate inference yields xy as a conclusion, then the counterexemplary extension seems to call into question the adequacy of the directly skeptical construction, more readily than inversely.

2.2 Upward vs. Downward Chaining

The view of paths as lines of reasoning leads naturally to inductive definitions of accepted paths: for $\sigma : x \rightarrow y \rightarrow \ldots \rightarrow z \mapsto u$ (where $\mapsto$ stands for $\rightarrow$ or $\not\rightarrow$) to be valid, either the initial segment $\sigma' : x \rightarrow y \rightarrow \ldots \rightarrow z$ or the end segment $\sigma" : y \rightarrow \ldots \rightarrow z \mapsto u$ has to be valid already. The first condition is called upward, the latter downward chaining, double chaining is the conjunction of both. At first sight, a decision between the two seems arbitrary. Yet diagram Γ_1 will show that downward chaining might violate the principle of specificity: Reasoning upward, we will accept $z \rightarrow u \rightarrow v \rightarrow y$, thus zy, because the possible preclusion $z \rightarrow u \rightarrow x \rightarrow v$ and $x \not\rightarrow y$ is itself destroyed by the still more specific information $z \not\rightarrow x$. In downward chaining, however, only properties of direct successors can be inherited. And in the example, no direct successor w of z has the accepted property wy, as the destruction $z \not\to x$ of $z \to u \to x \to v$ is no longer "visible" at u. It thus seems that in normal cases the result of upward chaining is the better one. Thus, there is some good reason to opt for upward chaining definitions. A word of warning: A superficial impression of [SL] might be that whilst upward chaining is tractable, downward chaining is not, thus giving an additional criterion in favour of upward chaining. A more careful study of their results reveals that they show that whilst finding extensions defined by upward chaining is tractable, finding extensions defined by *double* chaining is not. Certain problems in the extensions approach have led [T] and others to consider double chaining. This will be discussed in more detail in Part 4. From now on, all definitions considered shall be (at least) upward chaining.

2.3 On-Path vs. Off-Path Preclusion

This is a rather technical distinction, discussed in [THT]. Briefly, a path σ: $x \to \ldots \to y \to \ldots \to z$ and a direct link $y \not\to u$ is an off-path preclusion of τ: $x \to \ldots \to z \to \ldots \to u$, but an on-path preclusion only iff all nodes of τ between x and z lie *on the path* σ. Thus, e.g. Γ_1 shows only on-path preclusion. A second warning: The wording of the definition seems to be a little misleading. A precise definition of on-path preclusion is given implicitely in [T]: [THT] refers to its construction as being on-path.

2.4 Split-Validity vs. Total-Validity Preclusion

Consider again a preclusion $\sigma : u \to \ldots \to x \to \ldots \to v$, and $x \not\to y$ of $\tau : u \to \ldots \to v \to \ldots \to y$. Most definitions demand for the preclusion to be effective - i.e. to prevent τ from being accepted - that the total path σ is valid. Some ([GV], [KK], [KKW1], [KKW2], [LS]) content themselves with the combinatorially simpler separate (split) validity of the lower and upper parts of σ: $\sigma' : u \to \ldots \to x$ and $\sigma" : x \to \ldots \to v$. In Γ_2, it is easily seen that $\sigma : u \to x \to w \to v$, $x \not\to y$ is only a split valid preclusion, as the link $u \not\to w$ destroys σ as a whole. Thus, split validity preclusion will give here the definite result $u\overline{y}$. With total validity preclusion, the diagram has essentially the form of a Nixon Diamond. A supporting argument for total validity preclusion can be given as follows : If we do not accept σ as true, but only σ' and σ", the truth of σ" might fail to take into account the pecularities of u, for the part of x containing u might behave irregularly. For illustration, interpret Γ_2 by assigning subsets of the real line to objects, and (probabilistic) set-inclusion to arrows: u:=[-1,0], x:= [-10,100], w:=[0,1000], v:=[-1,10000], $y := [-0.5, 0] \cup [100, 10000]$. (A general problem with probabilistic interpretations of defeasible inheritance nets is discussed in [Sc1].) Using techniques as in Γ_1, one may have valid total preclusion, but invalid split preclusion too.

2.5 Intersection of Extensions vs. the Intersection of their Conclusion Sets

Going into more technical details now, we need more terminology. Let us call all sequences of concatenated arrows of a net Γ, positive or negative, generalized paths. Thus, valid paths are potential ones, and the latter are generalized paths. If x, y are nodes in Γ, [x,y] shall denote the minimal subgraph of Γ containing all generalized paths in Γ beginning in x and ending in y.

The problem is perhaps best illustrated by discussing an example, Γ_3. At the heart of Γ_3 is a Nixon-Diamond [a,g]. The diagram has two extensions. If we decide [a,g] positively, i.e. make $a \to d \to g$ valid, then $a \to d \to g \to f$ and $g \not\to h$ is a valid preclusion of $a \to f \to h$, so $a \to f \to h \to i$ and $h \not\to k$ is no valid preclusion of $a \to i \to k$. But $a \to d \to g \to l \to i$ together with $l \not\to k$ is then a valid one. A negative choice in [a,g] will make $a \to f \to h \to i$ valid (thus, with $h \not\to k$ a valid preclusion of $a \to i \to k$) and $a \to d \to g \to l \to i$ invalid. The reader should note that all preclusions are on-path: a simpler diagram with off-path preclusions can be found in [MS]. Thus, we have in all extensions a valid preclusion of $a \to i \to k$, a valid negative path from a to k, thus $a\overline{k}$, but neither a common valid preclusion of $a \to i \to k$, nor a common valid negative path from a to k. This diagram teaches us a number of things. 1) One might have "floating conclusions", i.e. a conclusion might hold in all extensions without a *common* valid path to support it. (A similar result can be found in [St], working with an array of Nixon Diamonds, instead of preclusions.) 2) Consequently, validity in the intersection of extensions should not be defined by the set of common paths, but by the set of common conclusions. 3) How on earth is a directly skeptical approach supposed to "remember" which decision and diamond its paths come from ? Indeed, it can be shown (see [Sc2]) that any direct definition of inheritance, which has a (fixed) finite number of truth values, and satisfies a very weak and natural but somewhat technical condition will fail to give the same results as the intersection of extensions. The details are too involved to be presented here. The result will be motivated and taken up again in section 3.1. 4) It reveals the intricate interplay between decisions (in Nixon Diamonds) and preclusions possible in extensions - see section 3.2.

3 PROBLEMS SPECIFIC TO CERTAIN APPROACHES

3.1 Discussion of the [HTT] Approach

First, a problem of "right conclusion, wrong argument" of [HTT]. Consider $\Gamma' := [u, y] \subseteq \Gamma_4$. [HTT]'s definition will make the path $u \to v \to y$ valid, despite the valid preclusion $u \to w \to v$, $w \not\to y$. The conclusion uy itself is, of course, correct by the validity of $u \to x \to y$, but the argument $u \to v \to y$ seems wrong. (A similar problems arises for Nute's system [NBD].) The distinction may seem petty, but look at the full diagram. [HTT] will accept the preclusion $u \to v \to y$ and $v \not\to z$ of $u \to x \to y \to z$ and thus arrive at $u\overline{z}$, whereas the approach "correct in argument" will see $u \to x \to y$ as the only valid path from u to y, so we have no preclusion, but only a Nixon Diamond situation: $u \to v \not\to z$ vs. $u \to x \to y \to z$, so no conclusion in the skeptical definition. Of course, there is an easy way out: Accept as valid only paths, that are not precluded by valid paths.

A much more interesting problem shows up in the "Double Diamond" $\Gamma' := [a, j] \subseteq \Gamma_5$. As pointed out already in [HTT], their definition gives a wrong result: By skepticism, there is no path from a to i, so, by upwards chaining, the potential path $a \to f \to g \to j$ is unchallenged. On the other hand, there is a genuine possibility for $a\overline{j}$, by the path $a \to f \to i \not\to j$, which shows up in one extension. (Again, the same problem can be found in [NBD].) Thus, by being skeptical in the small diagram [a,i], we end up credulous in [a,j] wrt. aj, as we suddenly believe in one choice, when the opposite is (almost) as good. This phenomenon, as well as the one in the full diagram Γ_4, may be seen as special cases of the following general phenomenon: *Usually*, inheritance constructions are rich enough to turn any difference between two definitions of validity $\models$ and $\models'$ either way round, so there is no proper inclusion in the sense that $\Gamma \models \sigma$ implies $\Gamma \models' \sigma$ but not vice versa. Of course, examples like $\models := \emptyset$ provide exceptions, but the statement is correct in most interesting cases. What can we do ? First, one might try to work with an alternation of credulous and skeptical standards: If we want to be skeptical wrt. aj, we have to be credulous wrt. ai, and somehow admit the possibility $a \rightarrow f \rightarrow i$. Apart from technical difficulties, which can be guessed in the full diagram Γ_5, where we would have to start skeptically in [a,l], and liberally in [a,i] to arrive at skepticism wrt. ap, this involves an element of "cheating": We have to look upward to the result we would like to have, at least counting the number of alternations of standards, and cannot proceed purely inductively any more. As a second possible solution, one may preserve the destructive capabilities of invalid paths up the inductive hierarchy, thus essentially introducing a third truth value (valid, invalid, invalid but still destructive). Indeed, there are solutions of the Double Diamond problem still in the spirit of a directly skeptical approach, which follow this strategy, either explicitely [Sc1], or implicitely [GV]. These will, however, fail to give intuitive results in other diagrams - measured by the existence of reasonable extensions. As mentioned in section 2.5, this failure has a deeper reason: under some weak assumptions, we can show that *in principle* no direct skeptical definition with finitely many truth values can match exactly the intersection of extensions. Moreover, complexity results ([SL] for on-path preclusion, [Sc1] for off-path, and slightly modified, for on-path preclusion) tell us that finding the intersection of extensions is NP-hard.

3.2 Discussion of Stein's Approach

This approach can be very roughly described as follows (see [St]): First, credulous extensions of Γ are defined - call them L-extensions, as they differ from all extensions we have considered: So far, specificity is left totally out of consideration. (Thus, e.g. Γ_1 will have an L-extension containing $u \rightarrow v \rightarrow y$.) In a second step, a relation of preference is defined between L- extensions.

Third, the intersection of preferred L-extensions is taken. A tractable algorithm is given, though without proof of equivalence to the definition. The decisive step is, of course, the second one. An L-extension X is said to be preferred over Y, iff Y supports a precluded path, which is not redundant, and not supported by X. The notion of preclusion used in [St] is (prima facie) a very simple one, much stricter than even on-path preclusion. The interesting condition is "redundancy". A redundant path contains redundant links, which are shortcuts of longer paths, to which there is no opposite alternative of equal strength. And here lies the crux of the matter! In Γ_3 e.g., $a \to f$ would be redundant in the absence of $a \to e \not\to g$. The mere negative *possibility* $a \to e \not\to g$ however makes $a \to f$ non-redundant. The end-result is that in Γ_3, the preference relation between L-extensions is empty, so that we still have no conclusion for $ak/a\overline{k}$, contradicting the "true" intersection of extensions. In a very rough summary, this approach first takes a superficial look at Nixon Diamonds, considers only the negative possibility, looks then at preclusions, and reconsiders (preferred) extensions again in the end. For details, we have to refer the reader to the original article.

4 THE EXTENSIONS APPROACH - COHERENCE PROPERTIES

In section 1, we have described extensions as reasonable maximal consistent subsets of the potential paths - where "reasonable" stood for taking specificity into account. There is a property subtler than consistency and specificity, which we might call coherence. Look at Γ_6. In upwards chaining definitions, there is nothing so far to prevent extensions containing $a \to u \to v \to y$, and $b \to u \to x \not\to y$, a phenomenon, called "capriciousness" by Thomason, we may call it horizontal incoherence. Here, we are mainly concerned with a more disturbing situation, vertical incoherence or "decoupling" (see [T]): Consider $[a, y] \subseteq \Gamma_6$. We might have a "strange" extension with $a \to u \to v \to y$ and $u \to x \not\to y$ as valid paths. A solution to both problems would be double chaining. Yet, as we have seen, this has undesirable consequences too, it is too radical a remedy. (The neglect of specificity in downwards chaining, discussed in section 2.2, applies a fortiori to double chaining). Looking back at Γ_1, we see that we had there a *good reason* (the preclusion $z \not\to x$) to conclude zy and $u\overline{y}$, whereas in the present case, the decoupling is capricious or "unforced". In other words, what *are* good reasons to accept a decoupling situation ? One reason evidently is preclusion as in Γ_1. The task of finding a systematic and satisfying definition of acceptable decoupling situations turns out to be quite complex. The authors have experimented with a few simple definitions, which all produced counterintuitive results. This can be illustrated by examples like (the admittedly pathological) Γ_7: We shall assume that the paths $b \rightarrow g \not\rightarrow o$, $c \rightarrow i \not\rightarrow o$, $d \rightarrow k \rightarrow o$, $e \rightarrow m \rightarrow o$ are already admitted to the extension. Let's look at the left hand side of the diagram. All negative paths $a \rightarrow \ldots \not\rightarrow o$ are impossible (off-path precluded). The positive paths $a \rightarrow \ldots \rightarrow o$ are in conflictual decoupling pairs (cdp) with $b \rightarrow g \not\rightarrow o$ or $c \rightarrow i \not\rightarrow o$. There is no positive support for the cdp's on the left. The situation on the right-hand side is just the opposite: all positive paths $a \rightarrow \ldots \rightarrow o$ are precluded, all negative ones in cdp's. Shall we admit any path from a to o ? Our (informal) answer to the conflictual decoupling problem will be, that conflictual decoupling pairs shall be admitted to an extension, provided the situation satisfies the following conditions. Again, the exact definitions are too complex to be presented and discussed here, the reader is referred to [Sc1] for details. Consider an extension E of some net Γ, and a cdp $<\tau,\rho>$, e.g. $\tau : x \rightarrow y \rightarrow \ldots \rightarrow u \rightarrow \ldots \rightarrow z$, $\rho : y \rightarrow \ldots \rightarrow v \rightarrow \ldots \not\rightarrow z$, where ρ is already admitted: $\rho \in E$. Let τ/ρ denote the potential path $x \rightarrow y \rightarrow \ldots \rightarrow v \rightarrow \ldots \not\rightarrow z$. (C1): The presence of a valid path τ' in E, parallel to τ, with same conclusion xz, should be considered good positive reason to choose τ. This is intuitively acceptable (but may violate the "correct argument" principle of section 3.1 above), and otherwise, stability will fail very badly. We have, however, to take care, that no two cdp's support each other, without any further independent support. (C2): The preclusion of all μ in E contradicting τ should be considered good negative reason to choose τ. (C3): If an initial segment of τ/ρ is not in E, this should be considered good negative reason to choose τ - see Γ_1. Any cdp $<\tau,\rho>$ not supported by reasons (C1), (C2) or (C3) should be considered shaky. This has two consequences: (C4): If an unforced cdp $<\tau,\rho>$ is in conflict with another potential path σ, which is not in a cdp, then σ should be chosen, as σ is vertically less contested than τ. (C5): If - in the absence of any σ as discussed in (C4) - two unforced cdp's $<\tau,\rho>$ and $<\tau',\rho'>$ are in conflict with each other, as in Γ_7, we can choose one arbitrarily, just as in the case of a Nixon-Diamond.

5 Acknowledgements

We would like to thank the authors of the cited articles for their patient explanations.

References

[GV] H.Geffner, T.Verma : Inheritance = Chaining+Defeat, UCLA Technical Report CSD-890039, R-129-L, June 1989, or : in Methodologies for Intelligent Systems, IV, ISMIS 89, Charlotte NC. 1989

[H] J.Horty : Some Direct Theories of Nonmonotonic Inheritance, To appear in : Inheritance Hierarchies in Knowledge Representation, M.Lenzerini, D.Nardi, M.Simi Eds.

[HTT] J.F.Horty, R.H.Thomason, D.S.Touretzky : A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks, Dept. Comp. Sci., Carnegie Mellon Univ., CMU-CS-87-175, October 1987 or : Proceedings AAAI-87 (1987), 358-363.

[KK] T.Krishnaprasad, M.Kifer : An Evidence-based Framework for a Theory of Inheritance, IJCAI 89

[KKW1] T.Krishnaprasad, M.Kifer, D.Warren : On the Circumscriptive Semantics of Inheritance Networks, ISMIS 89

[KKW2] T.Krishnaprasad, M.Kifer, D.Warren : On the Declarative Semantics of Inheritance Networks, IJCAI 89

[KS] H.Kautz, B.Selman : Hard Problems for Simple Default Logics, Proceedings First International Conference on Principles of Knowledge Representation and Reasoning, Toronto, Canada, 1989

[LS] F.Lin, Y.Shoham : Argument Systems : a uniform basis for nonmonotonic reasoning, Stanford Computer Science Dept. Report STAN-CS-89-1243, 1989

[MS] D.Makinson, K.Schlechta : On Some Difficulties in the Theory of Defeasible Inheritance Nets, Proceedings of the Tübingen Workshop on Semantic Networks and Nonmonotonic Reasoning, vol.2, SNS Bericht, Universität Tübingen, 1989

[NBD] D.Nute, D.Billington, K.De Coster : Defeasible Logic and Inheritance Hierarchies with Exceptions, IJCAI 1989 (submitted)

[Sa] E. Sandewall : Non-monotonic inference rules for multiple in- heritance with exceptions, Proceedings IEEE 74 (1986), 1345-1353

[Sc1] K. Schlechta : Defeasible Inheritance : Coherence Properties and Semantics, Proceedings of the Tübingen Workshop on Semantic Networks and Nonmonotonic Reasoning, vol.1, SNS Bericht, Universität Tübingen, 1989

[Sc2] K.Schlechta : Directly Skeptical Inheritance cannot Capture the Intersection of Extensions, Submitted to Journal of Logic and Computation

[SL] B.Selman, H.Levesque : The Tractability of Path-Based Inheritance, IJCAI 1989

[St] L.Stein : Skeptical Inheritance : Computing the Intersection of Credulous Extensions, IJCAI 1989

[T] D.Touretzky : The Mathematics of Inheritance Systems, Los Altos/ London, 1986

[THT] D.S.Touretzky, J.F.Horty, R.H.Thomason : A Clash of Intuitions : The Current State of Nonmonotonic Multiple Inheritance Systems, IJCAI 1987