



Remarks on Inheritance Systems

Karl Schlechta

► To cite this version:

| Karl Schlechta. Remarks on Inheritance Systems. 2007. hal-00117112v3

HAL Id: hal-00117112

<https://hal.science/hal-00117112v3>

Preprint submitted on 9 Mar 2007 (v3), last revised 19 Aug 2007 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

REMARKS ON INHERITANCE SYSTEMS

Karl Schlechta *

Laboratoire d'Informatique Fondamentale de Marseille †

March 9, 2007

Abstract

We try a conceptual analysis of inheritance diagrams, first in abstract terms, and then compare this analysis to “normality” and the “small/big sets” of preferential and related reasoning. The main ideas are about nodes as truth values and information sources, truth comparison by paths, accessibility or relevance of information by paths, relative normality and size, and prototypical reasoning.

We will also see that some of the major distinctions between inheritance formalisms are consequences of deeper and more general problems of treating conflicting information.

AMS Classification: 68T27, 68T30

Contents

1	INTRODUCTION	2
2	BASIC INGREDIENTS	5
3	DIRECTLY SCEPTICAL SPLIT VALIDITY UPWARD CHAINING OFF-PATH INHERITANCE	12
4	A DISCUSSION OF NORMALITY	16

*ks@cmi.univ-mrs.fr, karl.schlechta@web.de, <http://www.cmi.univ-mrs.fr/> ~ ks

†UMR 6166, CNRS and Université de Provence, Address: CMI, 39, rue Joliot-Curie, F-13453 Marseille Cedex 13, France

5	SMALL SETS	18
5.1	Information transfer by relative size of subsets	20
6	INHERITANCE AS A DISTANCE BASED THEORY OF PROTO-TYPES	29
7	ACKNOWLEDGEMENTS	32

1 INTRODUCTION

(Defeasible or nonmonotonic) Inheritance networks or diagrams:

Nonmonotonic Inheritance Systems or Networks describe situations like “normally, birds fly”, written $bird \rightarrow fly$. Exceptions are permitted. Abstractly, such systems are *DAG's*, directed, acyclic graphs, with two types of arrows, \rightarrow and \nrightarrow . We will use Γ etc. for such graphs, and σ etc. for paths.

To simplify matters, we assume that for no two nodes $x, y \in \Gamma$ $x \rightarrow y$ and $x \nrightarrow y$ are both in Γ , intuitively, that Γ is free from (hard) contradictions. This restriction is inessential for our purposes. We admit, however, soft contradictions, and preclusion, which allows to solve soft contradictions.

Inheritance networks were introduced about 20 years ago (see e.g. [Tou84], [Tou86], [THT87]), and exist in a multitude of more or less differing formalisms, see e.g. [Sch97-2] for a brief discussion. There still does not seem to exist a satisfying semantics for these networks. The author’s own attempt [Sch90] is an a posteriori semantics, which cannot explain or criticise or decide between the different formalisms. On the other hand, inheritance networks seem a quite natural way of treating defeasible information, so an attempt to present a semantics (which then, necessarily, will have to discuss natural concepts used in treating defeasible information) seems worth while. This is what we do here. But we will also see that we have to take decisions which go beyond inheritance theory, thus, we should not expect such decisions to be settled by an inheritance semantics.

Truth values, normality, and subset size:

We discuss first an abstract semantics, which seems quite natural, in terms of information, information sources, accessibility of information sources, or relevance of the information from those sources, and comparison of information strength. We then discuss the resulting concept of normality, i.e. relative normality and not absolute normality, as present e.g. in preferential structures, and consequently in related systems like P or R (see e.g. [Sch04] for a presentation). We then look at the resulting properties of (again relative) concepts

of “small” and “big” subsets, and the (weak) laws which seem adequate.

These relativized versions of size have then a double interest: to give a semantics to inheritance, and, perhaps even more important, to indicate in which directions the overly ideal absolute notions of size have to be modified to get closer to “dirty” common sense reasoning.

Finally, we change perspective, and consider inheritance networks as a distance based theory of prototypes. The information at each node A gives the prototype for that node A , and nodes B for which there is a valid path from B to A , will inherit as much properties as possible from A , respecting specificity in case of conflict.

The text is organized as follows. We present in this order:

- (1) The main conceptual ideas (Section 2).
- (2) An inheritance formalism which corresponds; it is directly sceptical, upward chaining, split validity and off-path preclusion based. We also show correspondence to the discussion in Section 2 (Section 3).
- (3) A discussion of the concept of relativized normality (Section 4).
- (4) A translation of our approach into relativized subset size and a proof of correspondence to the inheritance formalism of Section 3 (Section 5).
- (5) An elaboration of the concept of prototype, and transfer of prototype information from one domain to another, using and comparing to the small subset concept. This is based on two different notions of distance, resulting in a distance based theory of prototypes.

The form of the text is (partly) somewhat unusual, as the text is based on a number of statements, which give our conceptual ideas. They seem plausible to us, but cannot be proved. Thus, they are suggestions how to consider inheritance systems. If the reader does not like them, he probably should stop reading and close the text. We accompany these statements with comments and discussion.

These conceptual ideas make split validity off-path preclusion, and upward chaining, a natural choice.

Choosing direct scepticism is a decision beyond the scope of this article, and we just make it. It is a general question how to treat contradictory and absent information, and if they are equivalent or not. (The fundamental difference between intersection of extensions and direct scepticism for defeasible inheritance was shown in [Sch93].)

This generates another problem, essentially that of the treatment of a mixture of contradictory and concordant information of multiple strengths or truth values. We bundle the decision of this problem with that for direct scepticism into a “plug-in” decision, which will be used in three approaches: the conceptual ideas, the inheritance algorithm, and the choice of the reference class for subset size. It is thus well encapsulated, and independent from the context. Of course, this is also a problem and decision transcending inheritance

diagrams.

The decomposition of the subset size based approach is, to a certain degree, arbitrary. It seems plausible to us, others might try a different decomposition. Our justification is that we tried to stay close to the standard treatment (see e.g. systems P and R) and bundle the necessary modifications in a few central points.

Decisions and their wider background:

We take now a closer look at several decisions made in this text. They all (to a lesser degree perhaps (1)) concern a wider subject than only inheritance networks. Thus, it is not surprising that there are different formalisms for solving such networks, deciding one way or the other. But this multitude is not the fault of inheritance theory, it is only a symptom of a deeper question. We will only mention the decisions here in overview for a clearer overall picture, they will be discussed in detail below, as they involve sometimes quite subtle questions.

(1) Upward chaining against downward or double chaining

(2.1) Off-path against on-path preclusion

(2.2) Split validity preclusion against total validity preclusion

(3) Direct scepticism against intersection of extensions

(4) Treatment of mixed contradiction and preclusion situations, no preclusion by paths of the same polarity

(1) This can also be seen as a difference in reasoning from cause to effect vs. backward, looking for causes for an effect. (A word of warning: There is a well-known article [SL89] from which a superficial reader might conclude that upward chaining is tractable, and downward chaining is not. A more careful reading reveals that, on the negative side, they only show that double chaining is not tractable.) We will adopt upward chaining in all our approaches.

(2.1) and (2.2) Both are consequences of our decision to see valid paths also as an absolute comparison of truth values, independent of reachability of information. This question of absoluteness transcends obviously inheritance networks. Our decision is, of course, again uniform for all our approaches.

(3) This point, too, is much more general than the problems of inheritance.

(4) concerns the treatment of truth values in more complicated situations, where we have a mixture of agreeing and contradictory information. Again, this goes by far beyond inheritance networks.

We will group (3) and (4) together in one general, “plug-in” decision, to be found in all approaches we discuss.

Varia:

The text is NOT self-contained, and familiarity with the basic concepts of inheritance systems and nonmonotonic logics in general is assumed. For a presentation, the reader might look into [Sch97-2] and [Sch04].

A.Bochman, Israel, has pointed out to author work by J.Barwise, D.Gabbay, C.Hartonas, [BGH95], on information flow. This has a superficial resemblance with the present pages. But, first, the BGH work is much deeper into logic, presenting sequent calculi, completeness results, etc. Second, our work is on non-monotonic logics, which BGH is not, and our main thrust is a conceptual analysis of inheritance networks, also working with multiple truth values. But the basic ideas are about similar situations.

It is an open question if our approach can be generalized to other situations, like Theory Revision, etc.

2 BASIC INGREDIENTS

Statement 2.1

Direct arrows (negative or positive) represent information, valid for their source. Thus, in a set reading, if there is an arrow $A \rightarrow B$ in the diagram, most elements of A will be in B , in short: “most A ’s are B ’s”.

Comment 2.1

Consequently, our reading covers also enriched diagrams, where arbitrary information can be “appended” to a node.

Statement 2.2

Nodes are information sources. If $A \rightarrow B$ is in the diagram, A is the source of the information “most A ’s are B ’s”.

Statement 2.3

Nodes are also truth values. They are the strength of the information whose source they are. See also Statement 2.5.

Statement 2.4

A valid, composed or atomic positive path σ from U to A makes the information of source A accessible to U . One might also say that A ’s information becomes relevant to U .

Comment 2.2

An alternative way to see a source of information is to see it as a reason to believe the information it gives. U needs a reason to believe s.t., i.e. a valid path from U to the source of the information, and also a reason to disbelieve, i.e. if U' is below U , and U' does NOT believe some information of A , then either it has stronger information to the contrary, or there is not a valid path to A any more (and neither to any other possible source of this information).

(“Reason”, a concept very important in this context, was introduced by A.Bochman into the discussion.)

Statement 2.5

A valid, composed or atomic positive path σ from A' to A allows to compare the strength of information source A' with that of A : A' is stronger than A . (In the set reading, this comparison is the result of specificity: more specific information is considered more reliable.) If there is no such valid path, we cannot resolve contradictions between information from A and A' .

Comment 2.3

Thus, nodes in an inheritance diagram are also information sources of different strengths or different truth values. (For simplicity, and to connect to other situations, I will speak now about truth values.) Valid positive paths allow to compare these truth values.

Thus, in a given node U , information from A is accessible iff there is a valid positive path from U to A , and if information from A' is also accessible, and there is a valid positive path from A' to A , then, in case of conflict, information from A' wins over that from A , as A' has a better truth value.

(In the usual drawings of preclusion diagrams, “access” corresponds to the vertical paths in the lower part of the diagram, and “comparison” to the horizontal ones. The upper part of the diagram represents just information. Tweety has access to penguins and birds, the horizontal link from penguin to bird compares the strengths, and the fly/not fly arrows are the information. See Diagram 2.1.)

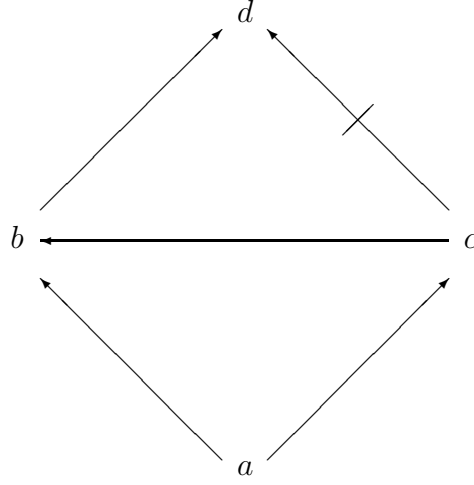
Comment 2.4

This reading justifies immediately the truth value order via valid paths as a specificity ordering. So it suffices to justify the specificity criterion.

We can do this as follows: If A is more specific than B , and contradicts B concerning ϕ , then there is a reason (A) to do so, but in all subsets of A this reason is by default valid. This is again due to a good choice of the language. Consider here the following artificial example, which serves to illustrate the importance of a well-chosen language.

The Tweety diagram

Read: a=Tweety, c=Penguins, b=Birds, d=Flying animals



In all diagrams, arrows point upwards, unless specified otherwise

Diagram 2.1

Example 2.1

Let the universe be a subset of the integers, $U := [1, \dots, 10] \cup [-11, \dots, -1]$, and define the predicates $P_1 := U - \{-11, -10\}$, $P_2 := P_1 - \{9, 10\}$, etc. Thus, in the universe the majority is negative, in P_1 positive, in P_2 negative again, etc., so P_n is better characterized by P_{n-2} than by P_{n-1} , and specificity gives the wrong answer - when we interpret “normal” by majority.

Comment 2.5

Thus, a negative direct link can only be information. A positive direct link is information at its source, but it can also be a comparison of truth values, or it can give access from its source to information at its end. A valid positive, composed path can only be comparison of truth values, or give access to information, it is NOT information itself.

(Note: This restriction applies to traditional inheritance networks, and the author makes no claim whatever that it should also hold for modified such systems, or in still other contexts. One of the reasons why we do not have “negative nodes”, and thus negated arrows also in the middle of paths might be the following (with \mathbf{C} complementation): If, for some X , we also have a node for $\mathbf{C}X$, then we should have $X \not\rightarrow \mathbf{C}X$ and $\mathbf{C}X \not\rightarrow X$,

thus a cycle, and arrows from Y to X should be accompanied by their opposite to $\mathbf{C}X$, etc. This would complicate the picture, probably without any real gain in insight.)

Comment 2.6

This interpretation results in split validity preclusion: the comparison between information sources A' and A is absolute, and does NOT depend on the U from which both may be accessible - as can be the case with total validity preclusion - see Diagram 6.9 in [Sch97-2] and its discussion there, which we repeat now, see Diagram 2.2 below. Here, the path $x \rightarrow w \rightarrow v$ is valid, so is $u \rightarrow x$, but not the whole preclusion path $u \rightarrow x \rightarrow w \rightarrow v$.

Split vs. total validity preclusion:

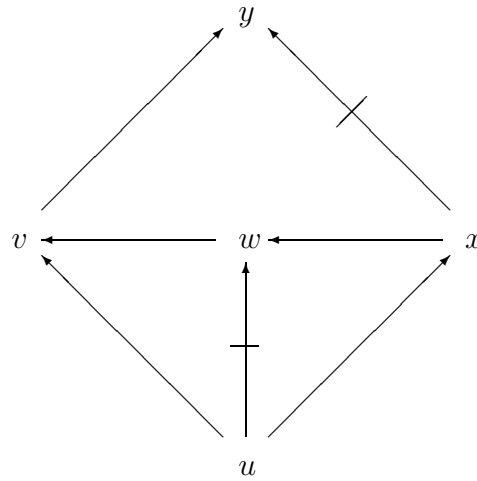


Diagram 2.2

Comment 2.7

Inheritance diagrams in this interpretation do not only represent reasoning with many truth values, but also reasoning ABOUT those truth values: their comparison is done by the same underlying mechanism.

Comment 2.8

We obtain automatically again that direct information is stronger than any other information: If A has information ϕ , and there is a valid path from A to B , making B 's information accessible to A , then this same path also compares strength, and A 's information is stronger than B 's information. Seen from A , i.e. just considering information accessible to A , A 's own information will always be best.

Comment 2.9

Our interpretation underlines the importance of initial segments: Initial segments make information accessible. Thus, initial segments have to be valid. Diagram 6.8, *p.* 179, in [Sch97-2] (which might be due to folklore of the field) shows requiring downward chaining would be wrong. We repeat it here, see Diagram 2.3.

The problem of downward chaining:

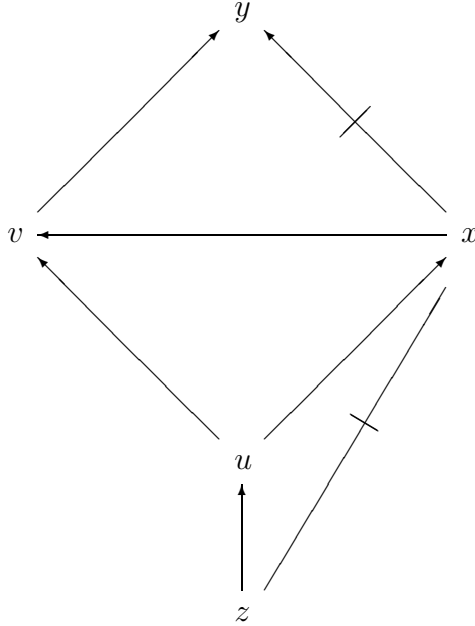


Diagram 2.3

Preclusions valid above (here at u) can be invalid at lower points (here at z), as part of the relevant information is not any more accessible (or becomes accessible). We have $u \rightarrow x \not\rightarrow y$ valid, by downward chaining, any valid path $z \rightarrow u \dots y$ has to have a valid final segment $u \dots y$, which can only be $u \rightarrow x \not\rightarrow y$, but intuition says that $z \rightarrow u \rightarrow v \rightarrow y$ should be valid. Downward chaining prevents such changes, and thus

seems inadequate, so we decide for upward chaining. (Already preclusion itself underlines upward chaining: In the Tweety diagram, we have to know that the path from bottom up to penguins is valid. So at least some initial subpaths have to be known - we need upward chaining.) (The rejection of downward chaining seems at first sight to be contrary to the intuitions carried by the word “inheritance”.)

Statement 2.6

Otherwise, information is considered independent from each other - only (valid) paths create the dependencies.

Definition 2.1

A plug-in decision:

We describe now a situation which we will meet in all contexts discussed, and whose decision goes beyond our problem - thus, we have to adopt one or several alternatives, and translate them into the approaches we will discuss. Thus, there must be one global decision, which is (and can be) adapted to the different contexts.

Suppose we have information about ϕ and ψ , where ϕ and ψ are presumed to be independent - in some adequate sense.

Suppose then that we have information sources $A_i : i \in I$ and $B_j : j \in J$, where the A_i speak about ϕ (they say ϕ or $\neg\phi$), and the B_j speak about ψ in the same way. Suppose further that we have a partial, not necessarily transitive (!), ordering $<$ on the information sources A_i and B_j together. $X < Y$ will say that X is better (intuition: more specific) than Y . (The potential lack of transitivity is crucial, as valid paths do not always concatenate to valid paths - just consider the tweety diagram.)

We also assume that there are contradictions, i.e. some A_i say ϕ , some $\neg\phi$, likewise for the B_j - otherwise, there are no problems in our context.

We can now take several approaches, all taking contradictions and the order $<$ into account.

- (P1) We use the global relation $<$, and throw away all information coming from sources of minor quality, i.e. if there is X s.t. $X < Y$, then no information coming from Y will be taken into account. Consequently, if Y is the only source of information about ϕ , then we will have no information about ϕ . This seems an overly radical approach, as one source might be better for ϕ , but not necessarily for ψ , too. If we adopt this approach, we can continue as below, and can even split in analogue ways into (P1.1) and (P1.2), as we do there for (P2.1) and (P2.2).
- (P2) We consider the information about ϕ separately from the information about ψ . Thus, we consider for ϕ only the A_i , for ψ only the B_j . Take now e.g. ϕ and the A_i . Again, there are (at least) two alternatives.

- (P2.1) We eliminate again all sources among the A_i for which there is a better $A_{i'}$, irrespective of whether they agree on ϕ or not. This is essentially the treatment of preclusion.
 - * (a) If the sources left are contradictory, we conclude nothing about ϕ , and accept for ϕ none of the sources. (This is a directly sceptical approach of treating unsolvable contradictions, following our general strategy.)
 - * (b) If the sources left agree for ϕ , i.e. all say ϕ , or all say $\neg\phi$, then we conclude ϕ (or $\neg\phi$), and accept for ϕ all the remaining sources.
- (P2.2) We eliminate again all sources among the A_i for which there is a better $A_{i'}$, but only if A_i and $A_{i'}$ have contradictory information. Thus, more sources may survive than in approach (P2.1).

We now continue as for (P2.1):

- * (a) If the sources left are contradictory, we conclude nothing about ϕ , and accept for ϕ none of the sources.
- * (b) If the sources left agree for ϕ , i.e. all say ϕ , or all say $\neg\phi$, then we conclude ϕ (or $\neg\phi$), and accept for ϕ all the remaining sources.

The difference between (P2.1) and (P2.2) is illustrated by the following simple example. Let $A < A' < A''$, but $A \not< A''$ (recall that $<$ is not necessarily transitive), and $A \models \phi$, $A' \models \neg\phi$, $A'' \models \neg\phi$. Then (P2.1) decides for ϕ (A is the only survivor), (P2.2) does not decide, as A and A'' are contradictory, and both survive in (P2.2).

There are arguments for and against either solution: (P2.1) gives a uniform picture, more independent from ϕ , (P2.2) gives more weight to independent sources, it “adds” information sources, and thus gives potentially more weight to information from several sources. (P2.2) seems more in the tradition of inheritance networks, so we will consider it in the further development.

Comment 2.10

We translate the analysis and decision of Definition 2.1 now into the picture of information sources, accessibility, and comparison via valid paths. This is straightforward:

- (1) We have that information from A_i , $i \in I$, about B is accessible from U , i.e. there are valid positive paths from U to all A_i . Some A_i may say $\neg B$, some B .
- (2) If information from A_i is comparable with information from A_j (i.e. there is a valid positive path from A_i to A_j or the other way around), and A_i contradicts A_j wrt. B , then the weaker information is discarded.
- (3) It remains a (nonempty, by lack of cycles) set of the A_i , s.t. for no such A_i there is A_j with contradictory information about B . If the information from this remaining set is contradictory, we accept none (and none of the paths either), if not, we accept the common conclusion and all these paths.

Statement 2.7

If we want to conform to inheritance, we must not add trivialities like “ x ’s are x ’s”, as this would require $x \rightarrow x$ in the corresponding net, which, of course, will not be there in an acyclic net.

3 DIRECTLY SCEPTICAL SPLIT VALIDITY UPWARD CHAINING OFF-PATH INHERITANCE

INHERITANCE)

Our approach is directly sceptical, i.e. unsolvable contradictions result in the absence of valid paths, it is upward chaining, and split-validity for preclusions. It is strongly inspired by classical work in the field by Horty, Thomason, Touretzky, and others, and we claim no priority whatever. If it is new at all, it is a very minor modification of existing formalisms.

We first define potential paths, then the notion of degree, and finally validity of paths, written $\Gamma \models \sigma$, if σ is a path.

The definition of \models (i.e. of validity of paths)

All definitions are relative to a fixed diagram Γ . The notion of degree will be defined relative to all nodes of Γ , as we will work with split validity preclusion, so the paths to consider may have different origins.

Definition 3.1

Potential paths:

If $x \rightarrow p \in \Gamma$, then $x \rightarrow p$ is a positive potential path (pp.).

If $x \not\rightarrow p \in \Gamma$, then $x \not\rightarrow p$ is a negative pp.

If $x \cdots \rightarrow p$ is a positive pp., and $p \rightarrow q \in \Gamma$, then $x \cdots \rightarrow p \rightarrow q$ is a positive pp.

If $x \cdots \rightarrow p$ is a positive pp., and $p \not\rightarrow q \in \Gamma$, then $x \cdots \rightarrow p \not\rightarrow q$ is a negative pp.

Definition 3.2

Degree:

As already indicated, we shall define paths inductively. As we do not admit cycles in our systems, the arrows define a well-founded relation on the vertices. Instead of using this relation for the induction, we shall first define the auxiliary notion of degree, and do induction on the degree. Given a node x (the origin), we need a mapping f from the

vertices to natural numbers s.t. $p \rightarrow q$ or $p \not\rightarrow q \in \Gamma$ implies $f(p) < f(q)$, and define (relative to x):

(1) a generalized path is any monotone coherent sequence of arrows of any type, beginning at x , like $x \rightarrow p \not\rightarrow q \not\rightarrow r \rightarrow s$ etc.

(2) Let σ be a generalized path from x to y , then $\deg_{\Gamma,x}(\sigma) := \deg_{\Gamma,x}(y) :=$ the maximal length of any generalized path parallel to σ , i.e. beginning in x and ending in y .

Definition 3.3

Inductive definition of $\Gamma \models \sigma$:

Let σ be a potential path.

- Case I:

σ is a direct link in Γ . Then $\Gamma \models \sigma$

(Recall that we have no hard contradictions in Γ .)

- Case II:

σ is a compound potential path, $\deg_{\Gamma}(\sigma) = n$, and $\Gamma \models \tau$ is defined for all τ with degree less than n - whatever their origin and endpoint.

- Case II.1:

Let σ be a positive pp. $x \cdots \rightarrow u \rightarrow y$, denote $\sigma' := x \cdots \rightarrow u$.

Then $\Gamma \models \sigma$ iff

- (1) (1) σ is a candidate by upward chaining,
- (2) (2) σ is not precluded by more specific contradicting information,
- (3) (3) all potential contradictions are themselves precluded by information contradicting them.

Note that (2) and (3) are the translation of (P2.2) in Definition 2.1.

More precisely:

- (1) (1) $\Gamma \models \sigma'$ and $u \rightarrow y \in \Gamma$.

(The initial segment must be a path, as we have an upward chaining approach. This is decided by induction hypothesis.)

- (2) (2) There are no v, τ, τ' s.t. $v \not\rightarrow y \in \Gamma$ and $\Gamma \models \tau := x \cdots \rightarrow v$ and $\Gamma \models \tau' := v \cdots \rightarrow u$.

(σ itself is not precluded by split validity preclusion and a contradictory link. Note that $\tau \circ v \not\rightarrow y$ (= the concatenation) need not be valid, it suffices that it is a better candidate (by τ' .)

- (3) (3) all potentially conflicting paths are precluded by information contradicting them:

For all v and τ such that $v \not\rightarrow y \in \Gamma$ and $\Gamma \models \tau := x \cdots \rightarrow v$ (i.e. for all potentially conflicting paths $\tau \circ v \not\rightarrow y$) there is z such that $z \rightarrow y \in \Gamma$ and either

$z = x$

(the potentially conflicting pp. is itself precluded by a direct link, which is thus also valid)

or

there are $\Gamma \models \rho := x \cdots \rightarrow z$ and $\Gamma \models \rho' := z \cdots \rightarrow v$ for suitable ρ and ρ' .

(Thus, there is a valid path from x to z , and z is more specific than v , so $\tau \circ v \not\rightarrow y$ is precluded. Again, $\rho \circ z \rightarrow y$ need not be a valid path, but it is a better candidate than $\tau \circ v \not\rightarrow y$ is, and as $\tau \circ v \not\rightarrow y$ is in simple contradiction, this suffices.

- Case II.2: The negative case is entirely symmetrical.

Definition 3.4

Finally, define $\Gamma \models xy$ iff there is $\sigma : x \rightarrow y$ s.th. $\Gamma \models \sigma$, likewise for $x\bar{y}$.

Diagram 3.1 shows the most complicated situation for the positive case.

We have to show now that above approach corresponds to the preceeding discussion.

Fact 3.1

Above definition and the one outlined in Definition 2.1 correspond.

Proof:

(Outline)

We argue for the result, the argument for valid paths is similar.

Consider then case (P2.2) in Definition 2.1, as well as Comment 2.10, and start from some x .

Case 1:

Direct links, $x \rightarrow z$ or $x \not\rightarrow z$.

By definition, as a direct link starts at x , the information z or $\neg z$ is stronger than all other accessible information. Thus, the link and the information will be valid in both approaches. Note that we assumed Γ free from hard contradictions.

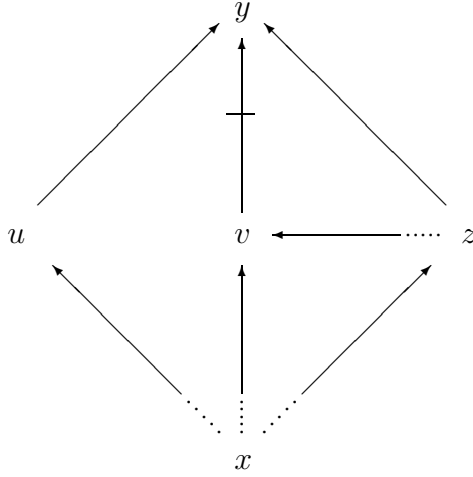


Diagram 3.1

Case 2:

Composite paths.

In both approaches, the initial segment has to be valid, as information will otherwise not be accessible. Also, in both approaches, information will have the form of direct links from the accessible source. Thus, condition (1) in Case II.1 corresponds to condition (1) in Comment 2.10.

In both approaches, information contradicted by a stronger source (preclusion) is discarded, as well as information which is contradicted by other, not precluded sources, so (P2.2) in Definition 2.1 and II.1 (2) + (3) correspond. Note that variant (P2.1) of Definition 2.1 would give a different result - which we could, of course, also imitate in a modified inheritance approach.

Case 3:

Other information.

Inheritance nets give no other information, and, as pointed out at the end of Comment 2.10, we do not add any other information either in the approach in Definition 2.1.

Thus, both approaches are equivalent.

□

4 A DISCUSSION OF NORMALITY

Statement 4.1

Normality in inheritance (and Reiter defaults etc.) is relative, and as much normality as possible is preserved. There is no $N(X)$, but only $N(X, \phi)$, and $N(X, \phi)$ might be defined, but not $N(X, \psi)$.

Comment 4.1

Normality in the sense of preferential structures is absolute: if x is not in $N(X)$ ($= \mu(X)$), we do not know anything beyond classical logic. This is the dark Swedes' problem: even dark Swedes should probably be tall. Inheritance systems are different: If birds usually lay eggs, then penguins, though abnormal wrt. flying, will still usually lay eggs. Penguins are fly-abnormal birds, but will continue to be egg-normal birds - unless we have again information to the contrary.

So the absolute, simple $N(X)$ of preferential structures splits up into many, by default independent, normalities, $N(X, \phi)$ for ϕ -normal etc. Preferential structures seem to be often an oversimplification.

This corresponds to intuition: There are no absolutely normal birds, each one is particular in some sense, so $\bigcap \{N(X, \phi) : \phi \in \mathcal{L}\}$ may well be empty, even if each single $N(X, \phi)$ is almost all birds. (Compare to the problem to define "human being".)

Comment 4.2

We try to preserve as much normality as possible. This can be pictured as follows: each ϕ -normality is a kind of wall, and normal elements are inside. If an element reveals itself as ϕ -abnormal, then we push it out of the ϕ -normality-wall, but keep it inside the other normalities, if possible. In addition, we cancel any properties which are based essentially on ϕ -normality.

So normality is treated by inertia: we give up only as little as necessary. An exception is: If the reason A to believe ϕ and ψ is not any more accessible, then we give up all information appended at A - but not more.

Comment 4.3

Both properties are based on default independence: by default, we presume properties to be independent, but if they are based on the same reason, then they all depend on the reason being accessible (if there is no better information to the contrary accessible).

Comment 4.4

This results in a minimal distance semantics: the universally normal situation is the prototype, and we try to stay as close as possible. If the situation is not prototypical, we do not give up all, but only as far as needed, which is the Hamming distance. These ideas are elaborated further in Section 6.

The reason is an economical one: we can stay (in thinking about it) close to the prototype, and its known properties, and an ontological one: our language was created to give by default independence of properties. If this independence does not exist, then we have a reason for it, and we know it. (This should still be further elaborated.)

There might be misunderstandings about the use of the word “distance” here. The author is fully aware that inheritance networks cannot be captured by distance semantics in the sense of preferential structures. But we do NOT think here of distances from one fixed ideal point, but of relativized distances: Every prototype is the origin of measurements. E.g., the bird prototype is defined by “flying, laying eggs, having feathers . . .”. So we presume that all birds have these properties of the prototype, i.e. distance 0 from the prototype. When we see that penguins do not fly, we move as little as possible from the bird prototype, so we give up “flying”, but not the rest. Thus, penguins (better: the penguin prototype) will have distance 1 from the bird prototype (just one property has changed). So there is a new prototype for penguins, and considering penguins, we will not measure from the bird prototype, but from the penguin prototype, so the point of reference changes. This is exactly as in distance semantics for theory revision, introduced in [LMS01], only the point of reference is not the old theory T , but the old prototype, and the distance is a very special one, counting properties assumed to be independent. (The picture is a little bit more complicated, as the loss of one property (flying) may cause other modifications, but the simple picture suffices for this informal argument, left for further elaboration.)

Comment 4.5

What are the laws of relative normality? Following above discussion, $N(X, \phi)$ and $N(X, \psi)$ will be largely independent (except for trivial situations, where $\phi \leftrightarrow \psi$, ϕ is a tautology, etc.). (In softening independence, can we say that $N(X, \phi)$ and $N(X, \psi)$ are close, if ϕ and ψ are close?) Below, in Section 6, where we pursue the idea of prototypes more seriously, we will postulate that $N(X, \phi)$ and $N(X, \phi')$ should at least not have empty intersection.

Comment 4.6

Note that, by relative normality, $N(X, \phi)$ might be defined, and $N(X, \psi)$ not. Thus, if there is no arrow, or no path, between X and Y , then $N(X, Y)$ and $N(Y, X)$ - where X, Y are also properties - need not be defined. This will get rid of the unwanted connections found with absolute normalities: Penguins don't fly, birds do, penguins are a subclass of birds, so normal birds are not penguins. Yet, there is no arrow $Birds \not\rightarrow Penguins$ in the usual Tweety diagram. But there are NO normal birds in the absolute sense, only ϕ -normal birds. $N(bird, penguin)$ is simply not defined. This is like a new truth value, "undefined". See the discussion in Comment 5.1 below.

Comment 4.7

Up to now, we have not really used the relative normalities. We expressed doubts that this can be done in a strong way "horizontally", combining $N(X, \phi)$ and $N(X, \psi)$ etc., but there seems a natural way to work with them "vertically", combining $N(X, Y)$ and $N(Y, \phi)$ etc. in a kind of normality transfer. It seems that this is done better in the language of "small" and "big" subsets, but we have to stress that these notions are not absolute for given X , but also depend on ϕ, ψ etc. We turn to such sets.

5 SMALL SETS

Statement 5.1

The usual informal way of speaking about inheritance networks (plus other considerations) motivates an interpretation by sets and soft set inclusion - $A \rightarrow B$ means that "most A 's are B 's". Just as with normality, the "most" will have to be relativized, i.e. there is a B -normal part of A , and a B -abnormal one, and the first is B -bigger than the second - where "bigger" is relative to B , too. A further motivation for this set interpretation is the often evoked specificity argument for preclusion. Thus, we will now translate our remarks about normality into the language of big and small subsets.

Statement 5.2

Recall our remarks about relative normality. $N(X, \phi)$ is, a priori, independent of $N(X, \psi)$, and $N(X, \phi)$ might be defined, but not $N(X, \psi)$. Thus, we will have ϕ -big subsets of X , and ψ -big subsets, and the two are independent, may have empty intersection, only one may be defined, etc.

Comment 5.1

Consider now the system P (with CUM). Small sets are used in two conceptually very distinct ways: $\alpha \sim \beta$ iff the set of $\alpha \wedge \neg\beta$ -cases is a small subset (in the absolute

sense, there is just one system of big subsets of the α -cases) of the set of α -cases. The second use is in information transfer, used in CUM, or CM more precisely: if the set of $\alpha \wedge \neg\gamma$ -cases is a small subset of the set of α -cases, then $\alpha \vdash \beta$ carries over to $\alpha \wedge \gamma : \alpha \wedge \gamma \vdash \beta$. (See also the discussion in [Sch04], page 86, after Definition 2.3.6.) It is this transfer which we will consider here, and not things like AND, which connect different $N(X, \phi)$ for different ϕ .

Before we go into details, we will show that e.g. the system P is too strong to model inheritance systems, and that e.g. the system R is too weak for this purpose. Thus, preferential systems are really quite different from inheritance systems.

First, P is too strong: Consider the diagram $A \rightarrow B \rightarrow C$, $A \not\rightarrow C$. There is no arrow $B \not\rightarrow A$, and we will see that P forces one to be there. For this, we take the natural translation, i.e. $X \rightarrow Y$ will be “ $X \cap Y$ is a big subset of X ”, etc. We show that $A \cap B$ is a small subset of B , which we write $A \cap B < B$. $A \cap B = (A \cap B \cap C) \cup (A \cap B \cap \mathbf{C}C)$. $A \cap B \cap \mathbf{C}C \subseteq B \cap \mathbf{C}C < B$, the latter by $B \rightarrow C$, thus $A \cap B \cap \mathbf{C}C < B$, essentially by Right Weakening. Set now $X := A \cap B \cap C$. As $A \not\rightarrow C$, $X := A \cap B \cap C \subseteq A \cap C < A$, and by the same reasoning as above $X < A$. It remains to show $X < B$. We use now $A \rightarrow B$. As $A \cap \mathbf{C}B < A$, and $A \cap X < A$, by Cumulativity $X = A \cap X \cap B < A \cap B$, so essentially by OR $X = A \cap X \cap B < B$. Using the filter property, we see that $A \cap B < B$.

Second, even R is too weak: In the diagram $X \rightarrow Y \rightarrow Z$, we want to conclude that most of X is in Z , but as X might also be a small subset of Y , we cannot transfer the information “most Y ’s are in Z ” to X .

Statement 5.3

We have to distinguish direct information/arrows from inherited information/valid paths. In the language of big/small sets, it is easiest to do this by two types of big subsets: big ones and very big ones. We will denote the first big, the second BIG, and bIG will denote big or BIG.

In particular, we will have the implications $BIG \rightarrow big$ and $SMALL \rightarrow small$, so we have nested systems. Such systems were discussed in [Sch95-1], see also [Sch97-2].

Comment 5.2

In particular, this distinction seems to be necessary to prevent arbitrary concatenation of valid paths to valid paths, which would lead to contradictions. Consider e.g. $a \rightarrow b \rightarrow c \rightarrow d$, $a \rightarrow e \not\rightarrow d$, $e \rightarrow c$. Then concatenating $a \rightarrow b$ with $b \rightarrow c \rightarrow d$, both valid, would lead to a simple contradiction with $a \rightarrow e \not\rightarrow d$, and not preclusion, as it should be.

Statement 5.4

For the situation $X \rightarrow Y \rightarrow Z$, we will need to conclude that:

If $Y \cap Z$ is a Z -BIG subset of Y and $X \cap Y$ is a Y -big subset of X then $X \cap Z$ is a Z -big subset of X .

Comment 5.3

$Y \rightarrow Z$ expresses the direct information in this context, so $Y \cap Z$ has to be a Z -BIG subset of Y .

$X \rightarrow Y$ can be direct information, but it is used here as channel of information flow, in particular it might be a composite valid path, so in our context, $X \cap Y$ is a Y -big subset of X .

$X \cap Z$ is a Z -big subset of X : this can only be big, and not BIG, as we have a composite path.

Comment 5.4

The translation into big/small subsets and their modifications is now quite complicated: we seem to have to relativize, and we seem to need two types of big/small. This casts, of course, a doubt on the enterprise of translation. Future will tell if any of the ideas can be used in other contexts.

We call the procedure formulated in Statement 5.4 information transfer, and analyze it now.

5.1 Information transfer by relative size of subsets

We investigate this situation now in more detail, first without conflicts.

The way we cut the problem is not the only possible one. We were guided by the idea that we should stay close to usual argumentation about big/small sets, should proceed carefully, i.e. step by step, and should take a general approach.

Note that we start without any X -big subsets defined, so X is not even a X -big subset of itself.

(A) The simple case of two arrows, and no conflicts.

If information ϕ is appended at Y , and Y is accessible from X (and there is no better information about ϕ available), ϕ will be valid at X . For simplicity, suppose there is a direct positive link from X to Y , written sloppily $X \rightarrow Y \models \phi$.

In the big subset reading, we will interpret this as: $Y \wedge \phi$ is a ϕ -BIG subset of Y . It is important that this is now direct information, so we have “BIG” and not “big”.

We read now $X \rightarrow Y$ also as: $X \cap Y$ is an Y -big subset of X - this is the channel, so just “big”.

We want to conclude by transfer that $X \cap \phi$ is a ϕ -big subset of X .

We do this in two steps: First, we conclude that $X \cap Y \cap \phi$ is a ϕ -big subset of $X \cap Y$, and then, as $X \cap Y$ is an Y -big subset of X , $X \cap \phi$ itself is a ϕ -big subset of X . We do NOT conclude that $(X - Y) \cap \phi$ is a ϕ -big subset of $X - Y$, this is very important, as we want to preserve the reason of being ϕ -big subsets - and this goes via Y !

The transition from “BIG” to “big” should be at the first step, where we conclude that $X \cap Y \cap \phi$ is a ϕ -big (and not ϕ -BIG) subset of $X \cap Y$, as it is really here where things happen, i.e. transfer of information from Y to arbitrary subsets $X \cap Y$.

Now, for the two steps in a slightly modified notation, corresponding to the diagram $X \rightarrow Y \rightarrow Z$:

(Here and in what follows, we will be very cautious, and relativize all normalities. We could perhaps obtain our objective with a more daring approach, using absolute normality here and there. But this would be a purely technical trick (interesting in its own), and we look here more for a conceptual analysis, and, as long as we do not find good conceptual reasons why to be absolute here and not there, we will just be relative everywhere.)

(1) If $Y \cap Z$ is a Z -BIG subset of Y (by $Y \rightarrow Z$), and $X \cap Y$ is a Y -big subset of X (by $X \rightarrow Y$), then $X \cap Y \cap Z$ is a Z -big subset of $X \cap Y$.

(2) If $X \cap Y \cap Z$ is a Z -big subset of $X \cap Y$, and $X \cap Y$ is a Y -big subset of X (by $X \rightarrow Y$) again, then $X \cap Z$ is a Z -big subset of X , so $X \dots \rightarrow Z$.

Note that (1) is very different from CUM or even RM, as we do not say anything about X in comparison to Y : X need not be any big or medium size subset of Y .

Seen as strict rules, this will not work, as it is transitivity, and thus monotony: we have to admit exceptions, as there might just be a negative arrow $X \not\rightarrow Z$ in the diagram. We will discuss such situations below in (C), where we will modify our approach slightly, and obtain a clean analysis.

We try now to give justifications for the two (defeasible) rules. They will be philosophical and can certainly be contested and/or improved.

For (1):

We look at Y . By $X \rightarrow Y$, Y 's information is accessible at X , so, as Z -BIG is defined for Y , Z -big will be defined for $Y \cap X$. Moreover, there is a priori nothing which prevents X from being independent from Y , i.e. $Y \cap X$ to behave like Y wrt. Z - by default: of course, there could be a negative arrow $X \not\rightarrow Z$, which would prevent this.

Thus, as $Y \cap Z$ is a Z -BIG subset of Y , $Y \cap X \cap Z$ should be a Z -big subset of $Y \cap X$. By the same argument (independence), we should also conclude that $(Y - X) \cap Z$ is a Z -big subset of $Y - X$. The definition of Z -big for $Y - X$ seems, however, less clear.

To summarize, $Y \cap X$ and $Y - X$ behave by default wrt. Z as Y does, i.e. $Y \cap X \cap Z$

is a Z -big subset of $Y \cap X$ and $(Y - X) \cap Z$ is a Z -big subset of $Y - X$. The reasoning is downward, from supersets to subsets, and symmetrical to $Y \cap X$ and $Y - X$. If the default is violated, we need a reason for it.

This default is an assumption about the adequacy of the language. Things do not change wildly from one concept to another (or, better: from Y to $Y \wedge X$), they might change, but then we are told so - by a corresponding negative link in the case of diagrams. See also Example 2.1.

For (2):

By $X \rightarrow Y$, X and Y are related, and we assume that X behaves as $Y \cap X$ does wrt. Z . This is upward reasoning, from subset to superset and it is NOT symmetrical: There is no reason to suppose that $X - Y$ behaves the same way as X or $Y \cap X$ do wrt. Z , as the only reason for Z we have, Y , does not apply.

Note that, putting relativity aside (which can also be considered as being big/small in various, per default independent dimensions) this is close to the reasoning with absolutely big/small sets: $X \cap Y - (X \cap Y \cap Z)$ is small in $X \cap Y$, so a fortiori small in X , and $X - (X \cap Y)$ is small in X , so $(X - (X \cap Y)) \cup (X \cap Y - (X \cap Y \cap Z))$ is small in X by the filter property, so $X \cap Y \cap Z$ is big in X , so a fortiori $X \cap Z$ is big in X .

Thus, in summary, we conclude by default that,

(3) If $Y \cap Z$ is a Z -BIG subset of Y , and $X \cap Y$ is a Y -big subset of X , then $X \cap Z$ is a Z -big subset of X .

(This analysis might need elaboration in future work.)

(B) The case with longer valid paths, but without conflicts.

Treatment of longer paths: Suppose we have a valid composed path from X to Y , $X \dots \rightarrow Y$, and not any longer a direct link $X \rightarrow Y$. By induction (upward chaining!) we argue - use directly (3) - that $X \cap Y$ is a Y -big subset of X , and conclude by (3) again that $X \cap Z$ is a Z -big subset of X .

(C) Treatment of multiple and perhaps conflicting information.

Consider the following example, illustrated in Diagram 5.1 :

Example 5.1

$X \rightarrow Y \rightarrow Z$, $X \rightarrow Y' \not\rightarrow Z$, $X \rightarrow U$, $Y' \rightarrow Y$, $Y'' \rightarrow Z$.

We want to analyze the situation and argue that e.g. X is mostly not in Z , etc.

First, all arguments about X and Z go via the Y 's. The arrows from X to the Y 's, and from Y' to Y could also be valid paths. We look at information which concerns Z (thus U is not considered), and which is accessible (thus Y'' is not considered). (We can be slightly

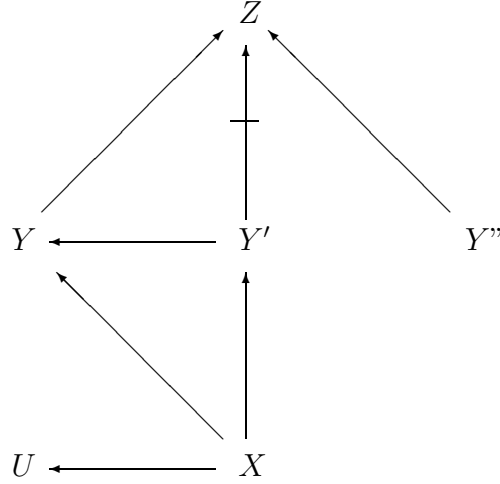


Diagram 5.1

more general, and consider all possible combinations of accessible information, not only those used in the diagram by X .) Instead of arguing on the level of X , we will argue one level above, on the Y 's and their intersections, respecting specificity and unresolved conflicts.

(Note that in more general situations, with arbitrary information appended, the problem is more complicated, as we have to check which information is relevant for some ϕ - conclusions can be arrived at by complicated means, just as in ordinary logic. In such cases, it might be better first to look at all accessible information for a fixed X , than at the truth values and their relation, and calculate closure of the remaining information.)

We then have (using the obvious language: “most A 's are B 's” for “ $A \cap B$ is a big subset of A ”, and “MOST A 's are B 's” for “ $A \cap B$ is a BIG subset of A ”).

In Y , Y'' , and $Y \cap Y''$, we have that MOST cases are in Z . In Y' and $Y \cap Y'$, we have that MOST cases are not in Z (= are in $\mathbf{C}Z$). In $Y' \cap Y''$ and $Y \cap Y' \cap Y''$, we are UNDECIDED about Z .

Thus:

$Y \cap Z$ will be a Z -BIG subset of Y , $Y'' \cap Z$ will be a Z -BIG subset of Y'' , $Y \cap Y'' \cap Z$ will

be a Z -BIG subset of $Y \cap Y''$.

$Y' \cap \mathbf{CZ}$ will be a Z -BIG subset of Y' , $Y \cap Y' \cap \mathbf{CZ}$ will be a Z -BIG subset of $Y \cap Y'$.

$Y' \cap Y'' \cap Z$ will be a Z -MEDIUM subset of $Y' \cap Y''$, $Y \cap Y' \cap Y'' \cap Z$ will be a Z -MEDIUM subset of $Y \cap Y' \cap Y''$.

This is just simple arithmetic of truth values, using specificity and unresolved conflicts, and the non-monotonicity is pushed into the fact that subsets need not preserve the properties of supersets.

In more complicated situations, we implement e.g. the general principle (P2.2) from Definition 2.1, to calculate the truth values. This will use in our case specificity for conflict resolution, but it is an abstract procedure, based on an arbitrary relation $<$.

This will result in the “correct” truth value for the intersections, i.e. the one corresponding to the other approaches.

It remains to do two things: (C.1) We have to assure that X “sees” the correct information, i.e. the correct intersection, and, (C.2), that X “sees” the accepted Y ’s, i.e. those through which valid paths go, in order to construct not only the result, but also the correct paths.

(Note that by split validity preclusion, if there is valid path from A through B to C , $\sigma : A \cdots \rightarrow B$, $B \rightarrow C$, and $\sigma' : A \cdots \rightarrow B$ is another valid path from A to B , then $\sigma' \circ B \rightarrow C$ will also be a valid path. Proof: If not, then $\sigma' \circ B \rightarrow C$ is precluded, but the same preclusion will also preclude $\sigma \circ B \rightarrow C$ by split validity preclusion, or it is contradicted, and a similar argument applies again.)

(C.1) Finding and inheriting the correct information:

X has access to Z -information from Y and Y' , so we have to consider them. Most of X is in Y , most of X is in Y' , i.e. $X \cap Y$ is a Y -big subset of X , $X \cap Y'$ is a Y' -big subset of X , so $X \cap Y \cap Y'$ is a $Y \cap Y'$ -big subset of X , thus most of X is in $Y \cap Y'$.

We thus have Y , Y' , and $Y \cap Y'$ as possible reference classes, and use specificity to choose $Y \cap Y'$ as reference class. We do not know anything e.g. about $Y \cap Y' \cap Y''$, so this is not a possible reference class.

Thus, we use specificity twice, on the Y ’s-level (to decide that $Y \cap Y'$ is mostly not in Z), and on X ’s-level (the choice of the reference class), but this is good policy, as, after all, much of nonmonotonicity is about specificity.

We should emphasize that nonmonotonicity lies in the behaviour of the subsets (determined by truth values and comparisons thereof) and the choice of the reference class by specificity. But both are straightforward now and local procedures, using information already decided before. There is no complicated issue here like determining extensions etc.

We now use above argument, described in the simple case, but with more detail, speaking

in particular about the most specific reference class for information about Z , $Y \cap Y'$ in our example - this is used essentially in (1.4), where the “real” information transfer happens, and here we also go from BIG to big.

(1.1) By $X \rightarrow Y$ and $X \rightarrow Y'$ (and there are no other Z -relevant information sources), we have to consider $Y \cap Y'$ as reference class.

(1.2) $X \cap Y$ is a Y -big subset of X (by $X \rightarrow Y$) (it is even Y -BIG, but we are immediately more general to treat valid paths), $X \cap Y'$ is a Y' -big subset of X (by $X \rightarrow Y'$). So $X \cap Y \cap Y'$ is a $Y \cap Y'$ -big subset of X .

(1.3) $Y \cap Z$ is a Z -BIG subset of Y (by $Y \rightarrow Z$), $Y' \cap \mathbf{C}Z$ is a Z -BIG subset of Y' (by $Y' \not\rightarrow Z$), so by preclusion $Y \cap Y' \cap \mathbf{C}Z$ is a Z -BIG subset of $Y \cap Y'$.

(1.4) $Y \cap Y' \cap \mathbf{C}Z$ is a Z -BIG subset of $Y \cap Y'$, and $X \cap Y \cap Y'$ is a $Y \cap Y'$ -big subset of X , so $X \cap Y \cap Y' \cap \mathbf{C}Z$ is a Z -big subset of $X \cap Y \cap Y'$.

This cannot be a strict rule without the reference class, as it would then apply to $Y \cap Z$, too, leading to a contradiction.

(2) If $X \cap Y \cap Y' \cap \mathbf{C}Z$ is a Z -big subset of $X \cap Y \cap Y'$, and $X \cap Y \cap Y'$ is a $Y \cap Y'$ -big subset of X , so $X \cap \mathbf{C}Z$ is a Z -big subset of X .

We make this now more formal.

We define for all nodes X, Y two sets: $B(X, Y)$, and $b(X, Y)$, where $B(X, Y)$ is the set of Y -BIG subsets of X , and $b(X, Y)$ is the set of Y -big subsets of X . (To distinguish undefined from medium/MEDIUM-size, we will also have to define $M(X, Y)$ and $m(X, Y)$, but we omit this here for simplicity.)

The translations are then:

(1.2') $X \cap Y \in b(X, Y)$ and $X \cap Y' \in b(X, Y') \Rightarrow X \cap Y \cap Y' \in b(X, Y \cap Y')$

(1.3') $Y \cap Z \in B(Y, Z)$ and $Y' \cap \mathbf{C}Z \in B(Y', Z) \Rightarrow Y \cap Y' \cap \mathbf{C}Z \in B(Y \cap Y', Z)$ by preclusion

(1.4') $Y \cap Y' \cap \mathbf{C}Z \in B(Y \cap Y', Z)$ and $X \cap Y \cap Y' \in b(X, Y \cap Y') \Rightarrow X \cap Y \cap Y' \cap \mathbf{C}Z \in b(X \cap Y \cap Y', Z)$ as $Y \cap Y'$ is the most specific reference class

(2') $X \cap Y \cap Y' \cap \mathbf{C}Z \in b(X \cap Y \cap Y', Z)$ and $X \cap Y \cap Y' \in b(X, Y \cap Y') \Rightarrow X \cap \mathbf{C}Z \in b(X, Z)$.

Finally:

(3') $A \in B(X, Y) \rightarrow A \in b(X, Y)$ etc.

Note that we used, in addition to the set rules, preclusion, and the correct choice of the reference class.

(C.2) Finding the correct paths:

Idea:

(1) If we come to no conclusion, then no path is valid, this is trivial.

(2) If we have a decision:

(2.1) All contradictory paths are out: e.g. $Y \cap Z$ will be Z -big, but $Y \cap Y' \cap \mathbf{C}Z$ will be Z -big. So there is no valid path via Y .

(2.2) Thus, not all paths supporting the same conclusion are valid.

Consider the following example, illustrated in Diagram 5.2 :

Example 5.2

$X \rightarrow Y \rightarrow Z$, $X \rightarrow Y' \not\rightarrow Z$, $X \rightarrow Y'' \rightarrow Z$, $Y'' \rightarrow Y' \rightarrow Y$.

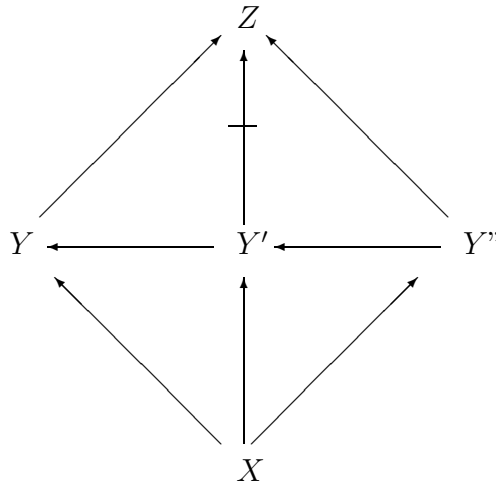


Diagram 5.2

There might be a positive path through Y , a negative one through Y' , a positive one through Y'' again, with $Y'' \rightarrow Y' \rightarrow Y$, so Y will be out, and only Y'' in. We can see this, as there is a subset, $\{Y, Y'\}$ which shows a change: $Y' \cap Z$ is Z -BIG, $Y' \cap \mathbf{C}Z$ is Z -BIG, $Y'' \cap Z$ is Z -BIG, and $Y \cap Y' \cap \mathbf{C}Z$ is Z -BIG, and the latter can only happen if there is a preclusion between Y' and Y , where Y loses. Thus, we can see this situation by looking only at the sets.

We show now equivalence with the inheritance formalism given in Section 3.

Fact 5.1

Above definition and the one outlined in Definition 3.1 correspond.

Proof:

By induction on the length of the deduction that $X \cap Z$ (or $X \cap \mathbf{C}Z$) is a Z -big subset of X . (Outline)

It is a corollary of the proof that we have to consider only subpaths and information of all generalized paths between X and Z .

Make all sets (i.e. one for every node) sufficiently different, i.e. all sets and boolean combinations of sets differ by infinitely many elements, e.g. $A \cap B \cap C$ will have infinitely many less elements than $A \cap B$, etc. (Infinite is far too many, we just choose it by laziness to have room for the $B(X, Y)$ and the $b(X, Y)$).

Put in $X \cap Y \in B(X, Y)$ for all $X \rightarrow Y$, and $X \cap \mathbf{C}Y \in B(X, Y)$ for all $X \not\rightarrow Y$ as base theory.

Length=1: Then big must be BIG, and, if $X \cap Z$ is a Z -BIG subset of X , then $X \rightarrow Z$, likewise for $X \cap \mathbf{C}Z$.

We stay close now to above example, so we argue for the negative case.

Suppose that we have deduced $X \cap \mathbf{C}Z \in b(X, Z)$, we show that there must be a valid negative path from X to Z . (The other direction is easy.)

Suppose for simplicity that there is no negative link from X to Z - otherwise we are finished.

As we can distinguish intersections from elementary sets (by the starting hypothesis about sizes), this can only be deduced using (2'). So there must be some suitable $\{Y_i : i \in I\}$ and we must have deduced $X \cap \bigcap Y_i \in b(X, \bigcap Y_i)$, the second hypothesis of (2'). If I is a singleton, then we have the induction hypothesis, so there is a valid path from X to Y . So suppose I is not a singleton. Then the deduction of $X \cap \bigcap Y_i \in b(X, \bigcap Y_i)$ can only be done by (1.2'), as this is the only rule having in the conclusion an elementary set on the left in $b(., .)$, and a true intersection on the right. Going back along (1.2'), we find $X \cap Y_i \in b(X, Y_i)$, and by induction hypothesis, there are valid paths from X to the Y_i .

The first hypothesis of (2') can be obtained by (1.3') or (1.4'). If it was obtained by (1.3'), then X is one of the Y_i , but then there is a direct link from X to Z (due to the "B", BIG). As a direct link always wins by specificity, the link must be negative, and we have a valid negative path from X to Z . If it was obtained by (1.4'), then its first hypothesis $\bigcap Y_i \cap \mathbf{C}B(\bigcap Y_i, Z)$ must have been deduced, which can only be by (1.3'), but the set of Y_i there was chosen to take all Y_i into account for which there is a valid path from X to

Y_i and arrows from the Y_i to Z (the rule was only present for the most specific reference class wrt. X and $Z!$), and we are done by the definition of valid paths in Section 3.

□

Comment 5.5

We summarize our ingredients.

Inheritance was done essentially by (1) and (2) of Section 5.1 (A) and its elaborations (1.i), (2) and (1.i'), (2'). It consisted of a mixture of bold and careful (in comparison to systems P and R) manipulation of big subsets. We had to be bolder than P/R are, as we have to transfer information also to perhaps small subsets. We had to be more careful, as P/R would have introduced far more connections than are present. We also saw that we are forced to loose the paradise of absolute small/big subsets, and have to work with relative size.

We then have a plug-in decision what to do with contradictions. This is a plug-in, as it is one (among many possible) solutions to the much more general question of how to deal with contradictory information, in the presence of a (partial, not necessarily transitive) relation which compares strength. At the same place of our procedure, we can plug in other solutions, so our approach is truly modular in this aspect. The comparing relation is defined by the existence of valid paths, i.e. by specificity.

This decision is inherited downward using again the specificity criterion.

Comment 5.6

Perhaps the deepest part of the analysis can be described as follows: Relevance is coded by positive arrows, and valid positive paths, and thus is similar to Kripke structures for modality. But, relevance (in this reading, which is closely related to causation) is profoundly non-monotonic, and any purely monotonic treatment of relevance would be insufficient. This seems to correspond to intuition. Relevance is then expressed formally by the possibility to combine different small/big sets.

This is, of course, a special form of relevance, there might be other forms of relevance.

Comment 5.7

In future work, one should probably isolate essentially two components:

- (1) Choosing the reference class
- (2) Inheriting the properties of the reference class to the class considered.

6 INHERITANCE AS A DISTANCE BASED THEORY OF PROTOTYPES

We change perspective slightly, and consider now information about classes, e.g. that birds fly, as the definition of (the properties of) a prototype, in this examples, of the bird prototype.

We preserve, however, our basic ideas to consider only direct links as information, and to use valid paths for truth value comparison, and for access to information.

We then have to address several questions:

- (1) What are essential and reasonable properties of prototypes?
- (2) How do prototypes relate e.g. to simple information sources?
- (3) How and under what circumstances is information passed from one prototype to another?
- (4) Is there any connection to the “small set approach” of Section 5?

We consider question (1).

In our simple setting, we can assume the properties of the prototype to be independent.

(1.1) In our simple case of independent information, we can describe the (information of the) prototype as a set of formulas.

(1.2) First, nothing is a prototype without information for what it is supposed to be a prototype. Thus, a prototype has a domain of, say, application. This gives, in the same time, a boundary for which it is defined as a prototype. Beyond its domain, it has no force any more. This does not mean that none of the properties it describes may hold elsewhere, but probably not all of them at the same place outside of its domain, and as we do not know which one might be preserved, it is reasonable not to say anything beyond the domain. In particular, if A is “almost” in B , and π is a prototype for B , then it has no information at all about $A - B$ (but perhaps about $A \cap B$). (This justifies also our complicated treatment of small sets, where we considered first $A \cap B$, and then went up to generalize to A .)

(1.3) In a nonmonotonic setting, a prototype π for domain A will admit exceptions, not all birds will fly, but most, etc. Thus, we can work with some sort of “big” subsets, i.e. every property will hold in some big subset.

(1.4) It might very well be that no element of the domain has ALL the properties of the prototype, every bird may be exceptional in some sense. It might, however, be reasonable to postulate that the big subset for property ϕ , and the one for property ϕ' have non-empty intersection. This will result in a “weak filter” and its axiomatization, as introduced by the author in [Sch95-1]. (Of course, it is trivial to generalize this concept to any finite

number n , i.e. any intersections of at most n ϕ_i -big subsets have non-empty intersection. The generalization is trivial on the semantic and the syntactic side.)

(1.5) Note that we might have situations like the Tweety diagram, where the only subset of birds is an abnormal one. A first idea might be that this is unnatural, most birds fly, and the only ones mentioned more specifically do not. But, we might just not have given the whole information, and not mentioned all the ravens, sparrows, tits, kolibris, etc. - so this is like a hierarchy still under construction, but not a contradiction or nonsense.

Concerning (2), (1.2) gives already a partial answer: An information source is not undefined, unless one knows its addressees, so it is a simpler concept.

Concerning (3), we remark:

(3.1) By (1.2), it does not seem reasonable to pass information of π to the outside of its domain.

(3.2) Let then B be a (strict) subset of the domain A of the prototype π . B can at the same time be a strict subset of the domain A' of the prototype π' . Let $\phi \in \pi$. If $\phi \in \pi'$, too, or $\neg\phi \notin \pi'$, it seems reasonable to pass ϕ to the prototype σ (to be constructed) of B . Thus, we have a maximality postulate, or minimal change change postulate: if there is no contradiction, we make σ as close to π as possible, where distance is measured by the Hamming distance.

(3.3) If, in the situation of (3.2), we have $\phi \in \pi$, and $\neg\phi \in \pi'$, we have a (local) contradiction, and it would be bad to take both information into σ . If we have no way to decide, we can branch into two extensions, or take none. The extensions approach presupposes that only the two cases are possible, yet this is not true in our context. It might well be that neither most elements of B satisfy ϕ , nor $\neg\phi$, as neither subset of B is big. Thus, tertium datur. So it seems better to take none.

In inheritance, we can decide by specificity. But, this can again be seen now as a distance based approach: If B is a (soft) subset of A and A' , and A is a (soft) subset of A' , then B is closer to or less different (as a soft subset) from A than it is to or from A' . In a qualitative reasoning, we would conclude that $A - B$ is smaller than $A' - B$, and as $B - A$ and $B - A'$ are small in B (and there is also no comparison between the two), the difference between B and A is smaller than the difference between B and A' .

Note that, again, automatically, B is closest to itself, so any direct link (=information) starting at B will win.

The difference to the distance based semantics for counterfactuals (or for theory revision) is that we do not choose the globally closest prototypes, but for each piece of information the closest prototype which has something to say about this information. It is as if we put out a question: does anyone know something about ϕ ? and take the answer from the closest prototype.

(3.4) Obviously, in other cases, it need not be specificity, which decides conflicts. Consider the case where Tweety is a bird, but a dead animal. Obviously, Tweety will not fly, here because the predicate “dead” is very strong and overrules many normal properties. When we generalize this, we might have a hierarchy of causes, where one overrules the other, or the result may be undecided. For instance, a falling object might be attracted in a magnetic field, but a gusty wind might prevent this, sometimes, with unpredictable results. This is then additional information (strength of cause), and this problem is not addressed directly in traditional inheritance networks, we would have to introduce a subclass “dead bird” - and subclasses often have properties of “pseudo-causes”, being a penguin probably is not a “cause” for not flying, nor bird for flying, still, things change from class to subclass for a reason.

(3.5) We make this a little more formal now.

(3.5.1) Our distance is NOT symmetric, as we do not inherit upwards.

(3.5.2) If two nodes B and A are not connected by a valid path from B to A , they have infinite distance. (Lack of symmetry is a trivial consequence now.)

(3.5.3) If the distances from B to A and from B to A' are finite, and comparable, as there is a valid path from A to A' , i.e. B is less distant from A than it is from A' , then information from A wins in case of conflict (otherwise we don't accept either information). The distance here is measured by (soft) set difference.

(3.5.4) We accept all information accessible by (3.5.3), and thus stay as close as possible to the (weighted set of) superset prototypes - distance measured by Hamming distance. Note that we use different distance measures in (3.5.3) and (3.5.4) - set difference and Hamming distance.

(3.5.5) To summarize: To construct a subset prototype, we take the closest prototypes - distance measured by set difference of the domains - and stay as close as possible to them by taking as much as possible of their information - distance measured by Hamming distance.

(4) One connection to the small set approach was already used: to compare distance of the domains by soft set difference. The other distance measure (Hamming distance) does not seem related to the concept of small sets. Finally, the transfer of information from one prototype to the other was described in the preceding sections by homogeneity, or homogenous distribution of small/big subsets, this picture does not change.

We should emphasize again that the present approach justifies in hindsight also our first downward, then upward transfer of information, as we cannot leave the domain of the prototype.

The distance approach might also be easier to generalize to other situations, e.g. where we have single elements, and not only sets.

7 ACKNOWLEDGEMENTS

Some of above reflections originated directly or indirectly from a long email discussion and criticisms with and from A.Bochman, Israel.

References

- [BGH95] J.Barwise, D.Gabbay, C.Hartonas: “On the Logic of Information Flow”, Journal of the IGPL, Vol. 3.1, 1995
- [HTT87] J.F.Horty, R.H.Thomason, D.S.Touretzky, “A Sceptical Theory of Inheritance in Nonmonotonic Semantic Networks”, Dept. Comp. Sci., Carnegie Mellon Univ., CMU-CS-87-175, October 1987 also in: Proceedings AAAI-87 (1987), p.358-363, or: Artificial Intelligence 42 (1990), p.311-348
- [LMS01] D.Lehmann, M.Magidor, K.Schlechta: “Distance Semantics for Belief Revision”, Journal of Symbolic Logic, Vol.66, No. 1, March 2001, *p.* 295-317
- [SL89] B.Selman, H.Levesque, “The Tractability of Path-Based Inheritance”, Proceed. IJCAI 1989, p.1140-1145
- [Sch04] K.Schlechta: “Coherent Systems”, Elsevier, Amsterdam, 2004
- [Sch90] K.Schlechta, “Semantics for Defeasible Inheritance”, in: L.G.Aiello (ed.), “Proceedings ECAI 90”, London, 1990, p.594-597
- [Sch93] K.Schlechta, “Directly Sceptical Inheritance cannot Capture the Intersection of Extensions”, Journal of Logic and Computation, Oxford, Vol.3, No.5, pp.455-467, 1995
- [Sch95-1] K.Schlechta: “Defaults as generalized quantifiers”, Journal of Logic and Computation, Oxford, Vol.5, No.4, p.473-494, 1995
- [Sch97-2] K.Schlechta: “Nonmonotonic logics - Basic Concepts, Results, and Techniques” Springer Lecture Notes series, LNAI 1187, Jan. 1997, 243pp
- [THT87] D.S.Touretzky, J.F.Horty, R.H.Thomason, “A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems”, Proceed. IJCAI 1987, p.476-482
- [Tou84] D.S.Touretzky, “Implicit Ordering of Defaults in Inheritance Systems”, Proceed. AAAI 84, p.322-325

- [Tou86] D.S.Touretzky, “The Mathematics of Inheritance Systems”, Los Altos/ London, 1986