



HAL
open science

Remarks on Inheritance Systems

Karl Schlechta

► **To cite this version:**

| Karl Schlechta. Remarks on Inheritance Systems. 2007. hal-00117112v4

HAL Id: hal-00117112

<https://hal.science/hal-00117112v4>

Preprint submitted on 19 Aug 2007

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.

AN ANALYSIS OF DEFEASIBLE INHERITANCE SYSTEMS

Karl Schlechta *

Laboratoire d'Informatique Fondamentale de Marseille †

August 18, 2007

Abstract

We give a conceptual analysis of (defeasible or nonmonotonic) inheritance diagrams, and compare our analysis to the "small/big sets" of preferential and related reasoning.

In our analysis, we consider nodes as information sources and truth values, direct links as information, and valid paths as information channels and comparisons of truth values. This results in an upward chaining, split validity, off-path preclusion inheritance formalism.

We show that the small/big sets of preferential reasoning have to be relativized if we want them to conform to inheritance theory, resulting in a more cautious approach, perhaps closer to actual human reasoning.

Finally, we interpret inheritance diagrams as theories of prototypical reasoning, based on two distances: set difference, and information difference.

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

*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

2	Introduction to nonmonotonic inheritance	4
2.1	Basic discussion	4
2.2	Directly sceptical split validity upward chaining off-path inheritance . . .	8
2.3	Review of other approaches and problems	17
3	Introduction to small and big sets and the logical systems P and R	19
4	Interpretations	24
4.1	Introduction	24
4.2	Informal comparison of inheritance with the systems P and R	24
4.3	Inheritance as information transfer	25
4.4	Inheritance as reasoning with prototypes	30
5	Detailed translation of inheritance to modified systems of small sets	34
5.1	Normality	34
5.2	Small sets	34
6	ACKNOWLEDGEMENTS	44

1 Introduction

Throughout this paper, "inheritance" will stand for "nonmonotonic or defeasible inheritance". We will use indiscriminately "inheritance system", "inheritance diagram", "inheritance network", "inheritance net".

In this introduction, we first give the motivation for this article, then describe in very brief terms the basic components of inheritance diagrams, mention the basic ideas of our analysis, as well as some more general decisions about treating contradictory information.

Motivation

Inheritance systems or diagrams have an intuitive appeal. They seem close to human reasoning, natural, and are also implemented (see [Mor98]). Yet, they are a more procedural approach to nonmonotonic reasoning, and, to the author's knowledge, a conceptual analysis, leading to a formal semantics, as well as a comparison to more logic based formalisms like the systems P and R of preferential systems are lacking. We attempt to reduce the gap between the more procedural and the more analytical approaches in this particular

case. This will also give indications how to modify systems P and R to approach them more to actual human reasoning. Moreover, we establish a link to multi-valued logics and the logics of information sources (see e.g. [ABK07] and forthcoming work of the same authors, and also [BGH95]).

Inheritance diagrams:

Inheritance diagrams are deceptively simple. Their conceptually complicated nature is seen by e.g. the fundamental difference between direct links and valid paths, and the multitude of existing formalisms, upward vs. downward chaining, intersection of extensions vs. direct scepticism, on-path vs. off-path preclusion (or pre-emption), split validity vs. total validity preclusion etc., to name a few. Such a proliferation of formalisms usually hints at deeper problems on the conceptual side, i.e. that the underlying ideas are ambiguous, and not sufficiently analysed. Therefore, any clarification and resulting reduction of possible formalisms seems a priori to make progress. Such clarification will involve conceptual decisions, which need not be shared by all, they can only be suggestions. Of course, a proof that such decisions are correct is impossible, and so is its contrary.

Our analysis:

We will introduce into the analysis of inheritance systems a number of concepts not usually found in the field, like multiple truth values, access to information, comparison of truth values, etc. We think that this additional conceptual burden pays off by a better comprehension and analysis of the problems behind the surface of inheritance.

We will also see that some distinctions between inheritance formalisms go far beyond questions of inheritance, and concern general problems of treating contradictory information - isolating some of these is another objective of this article.

The text is essentially self-contained, still some familiarity with the basic concepts of inheritance systems and nonmonotonic logics in general is helpful. For a presentation, the reader might look into [Sch97-2] and [Sch04].

A.Bochman, 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.

The text is organized as follows. After an introduction to inheritance theory and big/small subsets and the systems P and R in Section 2 and Section 3, we turn to an informal description of the fundamental differences between inheritance and the systems P and R in Section 4.2, give an analysis of inheritance systems in terms of information and

information flow in Section 4.3, then in terms of reasoning with prototypes in Section 4.4, and conclude in Section 5 with a translation of inheritance into (necessarily deeply modified) coherent systems of big/small sets, respectively logical systems P and R .

One of the main modifications will be to relativize the notions of small/big, which thus become less "mathematically pure" but perhaps closer to actual use in "dirty" common sense reasoning.

2 Introduction to nonmonotonic inheritance

2.1 Basic discussion

We give here an informal discussion. The reader unfamiliar with inheritance systems should consult in parallel Definition 2.3 and Definition 2.4. As there are many variants of the definitions, it seems reasonable to discuss them before a formal introduction, which, otherwise, would pretend to be definite without being so.

(Defeasible or nonmonotonic) inheritance networks or diagrams

Nonmonotonic inheritance systems describe situations like "normally, birds fly", written $birds \rightarrow fly$. Exceptions are permitted, "normally penguins *don't* fly", $penguins \not\rightarrow fly$.

Definition 2.1

(+++*** Orig. No.: Definition 2.1)

A nonmonotonic inheritance net is a finite DAG, directed, acyclic graph, with two types of arrows or links, \rightarrow and $\not\rightarrow$, and labelled nodes. We will use Γ etc. for such graphs, and σ etc. for paths - the latter to be defined below.

Roughly (and to be made precise and modified below, we try to give here just a first intuition), $X \rightarrow Y$ means that "normal" elements of X are in Y , and $X \not\rightarrow Y$ means that "normal" elements of X are not in Y . In a semi-quantitative set interpretation, we will read "most elements of X are in Y ", "most elements of X are not in Y ", etc. This is by no means the only interpretation, as we will see.

According to the set interpretation, we will also use informally expressions like $X \cap Y$, $X - Y$, $\mathbf{C}X$, etc. But we will also use nodes informally as formulas, like $X \wedge Y$, $X \wedge \neg Y$, $\neg X$, etc. All this will only be used as an appeal to intuition.

Nodes at the beginning of an arrow can also stand for individuals, so $Tweety \not\rightarrow fly$ means something like: "normally, Tweety will not fly". As always in nonmonotonic systems, exceptions are permitted, so the soft rules "birds fly", "penguins *don't* fly",

and (the hard rule) "penguins are birds" can coexist in one diagram, penguins are then abnormal birds (with respect to flying). The direct link *penguins* \nrightarrow *fly* will thus be accepted, or considered valid, but not the composite path *penguins* \rightarrow *birds* \rightarrow *fly*, by specificity - see below. This is illustrated by Diagram 2.1, where *a* stands for Tweety, *c* for penguins, *b* for birds, *d* for flying animals or objects.

(Remark: The arrows $a \rightarrow c$, $a \rightarrow b$, and $c \rightarrow b$ can also be composite paths - see below for the details.)

The Tweety diagram




Diagram 2.1

(Of course, there is an analogous case for the opposite polarity, i.e. when the arrow from *b* to *d* is negative, and the one from *c* to *d* is positive.)

The main problem is to define in an intuitively acceptable way a notion of valid path, i.e. concatenations of arrows satisfying certain properties.

We will write $\Gamma \models \sigma$, if σ is a valid path in the network Γ , and if x is the origin, and y the endpoint of σ , and σ is positive, we will write $\Gamma \models xy$, i.e. we will accept the conclusion that x 's are y 's, and analogously $\Gamma \models x\bar{y}$ for negative paths. Note that we will not accept any other conclusions, only those established by a valid path, so many questions about conclusions have a trivial negative answer: there is obviously no path from x to y . E.g., there is no path from b to c in Diagram 2.1. Likewise, there are no

disjunctions, conjunctions etc. in our conclusions, and negation is present only in a strong form: "it is not the case that x's are normally y's" is not a possible conclusion, only "x's are normally not y's" is one. Also, possible contradictions are contained, there is no EFQ.

To simplify matters, we assume that for no two nodes $x, y \in \Gamma$ $x \rightarrow y$ and $x \not\rightarrow y$ 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 us to solve some soft contradictions - as we already did in the penguins example. We will also assume that all arrows stand for rules with possibly exceptions, again, this restriction is not important for our purposes. Moreover, in the abstract treatment, we will assume that all nodes stand for sets, though this will not be true for all examples discussed.

This might be the place for a remark on absence of cycles. Suppose we also have a positive arrow from b to c in Diagram 2.1. Then, the concept of preclusion collapses, as there are now equivalent arguments to accept $a \rightarrow b \rightarrow d$ and $a \rightarrow c \not\rightarrow d$. Thus, if we do not want to introduce new complications, we cannot rely on preclusion to decide conflicts. It seems that this will change the whole outlook on such diagrams.

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. We will give here a conceptual analysis, which provides also at least some building blocks for a semantics, and a translation into (a modified version of) the language of small/big subsets, familiar from preferential structures, see below in Section 3.

We will now discuss the two fundamental situations of contradictions, then give a detailed inductive definition of valid paths for a certain formalism so the reader has firm ground under his feet, and then present briefly some alternative formalisms.

As in all of nonmonotonic reasoning, the interesting questions arise in the treatment of contradictions and exceptions. The difference in quality of information is expressed by "preclusion" (or "*pre-emption*"). The basic diagram is the Tweety diagram, see Diagram 2.1.

Unresolved contradictions give either rise to a branching into different extensions, which may roughly be seen as maximal consistent subsets, or to mutual cancellation in directly sceptical approaches. The basic diagram for the latter is the Nixon Diamond, see Diagram 2.2, where $a = Nixon$, $b = Quaker$, $c = Republican$, $d = pacifist$.

In the directly sceptical approach, we will not accept any path from a to d as valid, as there is an unresolvable contradiction between the two candidates.

The extensions approach can be turned into an indirectly sceptical one, by forming first all extensions, and then taking the intersection of either the sets of valid paths, or of valid

The Nixon Diamond




Diagram 2.2

conclusions, see [MS91] for a detailed discussion.

In more detail:

Preclusion

In the above example, our intuition tells us that it is not admissible to conclude from the fact that penguins are birds, and that most birds fly that most penguins fly. The horizontal arrow $c \rightarrow b$ together with $c \not\rightarrow d$ bars this conclusion, it expresses specificity. Consequently, we have to define the conditions under which two potential paths neutralize each other, and when one is victorious. The idea is as follows: 1) We want to be sceptical, in the sense that we do not believe every potential path. We will not arbitrarily chose one either. 2) Our scepticism will be restricted, in the sense that we will often make well defined choices for one path in the case of conflict: a) If a compound potential path is in conflict with a direct link, the direct link wins. b) Two conflicting paths of the same type neutralize each other, as in the Nixon Diamond, where neither potential path will be valid. c) More specific information will win over less specific one.

(It is essential in the Tweety diagram that the arrow $c \not\rightarrow d$ is a direct link, so it is in

a way stronger than compound paths.) The arrows $a \rightarrow b$, $a \rightarrow c$, $c \rightarrow b$ can also be composite paths: The path from c to b (read $c \subseteq \dots \subseteq b!$), however, tells us, that the information coming from c is more specific (and thus considered more reliable), so the negative path from a to d via c will win over the positive one via b . The precise inductive definition will be given below. This concept is evidently independent of the length of the paths, $a \dots \rightarrow c$ may be much longer than $a \dots \rightarrow b$, so this is not shortest path reasoning (which has some nasty drawbacks, discussed e.g. in [HTT87]).

Before we give a formalism based on these ideas, we refine them, adopt one possibility (but indicate some modifications), and discuss alternatives later.

2.2 Directly sceptical split validity upward chaining off-path inheritance

Our approach will be directly sceptical, i.e. unsolvable contradictions result in the absence of valid paths, it is upward chaining, and split-validity for preclusions (discussed below). We will indicate modifications to make it extension based, as well as for total-validity preclusion. This approach 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.

Our conceptual ideas to be presented in detail in Section 4.3 make split validity, off-path preclusion and upward chaining a natural choice.

For the reader's convenience, we give here a very short resume of these ideas: We consider only arrows as information, e.g. $a \rightarrow b$ will be considered information b valid at or for a . Valid (composed positive) paths will not be considered (direct) information. They will be seen as a way to obtain information, so a valid path $\sigma : x \dots \rightarrow a$ makes information b accessible to x , and, secondly, as a means of comparing information strength, so a valid path $\sigma : a \dots \rightarrow a'$ will make information at a stronger than information at a' . Valid negative paths have no function, we will only consider the positive initial part as discussed above, and the negative end arrow as information, but never the whole path.

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, see the remark in Section 4.4. (The fundamental difference between intersection of extensions and direct scepticism for defeasible inheritance was shown in [Sch93].)

We turn now to the announced variants as well as a finer distinction within the directly sceptical approach.

Our approach generates another problem, essentially that of the treatment of a mixture of contradictory and concordant information of multiple strengths or truth values. We bun-

dle 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 (and implicitly also for the treatment as a prototype theory). It is thus well encapsulated, and independent from the context.

These decisions (but, perhaps to a lesser degree, (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 first give an overview for a clearer overall picture, and discuss them 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 reasoning, 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, the authors only show that double chaining is not tractable.) We will adopt upward chaining in all our approaches. See Section 4.4 for more remarks.

(2.1) and (2.2) Both are consequences of our view - to be discussed below in Section 4.3 - 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. It is, among other things, a question of whether only the two possible cases (positive and negative) may hold, or whether there might be still other possibilities. See Section 4.4.

(4) This concerns the treatment of truth values in more complicated situations, where we have a mixture of agreeing and contradictory information. Again, this problem reaches 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.

Definition 2.2

(+++*** Orig. No.: Definition 2.2)

A plug-in decision:

We describe now more precisely 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. There will 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 such that $X < Y$, 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 below 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.
 - * (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.

The reader should note that our approach is quite far from a fixed point approach in two ways: First, fixed point approaches seem more appropriate for extensions based approaches, as both try to collect a maximal set of uncontradictory information. Second, we eliminate information when there is better, contradicting information, even if the final result agrees with the first. This, too, contradicts in spirit the fixed point approach.

After these preparations, we turn to a formal definition of validity of paths.

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. For simplicity, we consider Γ to be just a set of points and arrows, thus e.g. $x \rightarrow y \in \Gamma$ and $x \in \Gamma$ are defined, when x is a point in Γ , and $x \rightarrow y$ an arrow in Γ . Recall that we have two types of arrows, positive and negative ones.

We first define generalized and potential paths, then the notion of degree, and finally validity of paths, written $\Gamma \models \sigma$, if σ is a path, as well as $\Gamma \models xy$, if $\Gamma \models \sigma$ and $\sigma : x \dots \rightarrow y$.

Definition 2.3

(+++*** Orig. No.: Definition 2.3)

(1) Generalized paths:

A generalized path is an uninterrupted chain of positive or negative arrows pointing in the same direction, more precisely:

$x \rightarrow p \in \Gamma \rightarrow x \rightarrow p$ is a generalized path,

$x \not\rightarrow p \in \Gamma \rightarrow x \not\rightarrow p$ is a generalized path.

If $x \cdots \rightarrow p$ is a generalized path, and $p \rightarrow q \in \Gamma$, then $x \cdots \rightarrow p \rightarrow q$ is a generalized path,

if $x \cdots \rightarrow p$ is a generalized path, and $p \not\rightarrow q \in \Gamma$, then $x \cdots \rightarrow p \not\rightarrow q$ is a generalized path.

(2) Concatenation:

If σ and τ are two generalized paths, and the end point of σ is the same as the starting point of τ , then $\sigma \circ \tau$ is the concatenation of σ and τ .

(3) Potential paths (pp.):

A generalized path, which contains at most one negative arrow, and then at the end, is a potential path. If the last link is positive, it is a positive potential path, if not, a negative one.

(4) 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 (partial) mapping f from the vertices to natural numbers such that $p \rightarrow q$ or $p \not\rightarrow q \in \Gamma$ implies $f(p) < f(q)$, and define (relative to x) :

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 2.4

(+++*** Orig. No.: Definition 2.4)

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, $\text{deg}_{\Gamma,a}(\sigma) = n$, 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. $x \cdots \rightarrow u \rightarrow y$, let $\sigma' := x \cdots \rightarrow u$, so $\sigma = \sigma' \circ u \rightarrow y$

Then, intuitively, $\Gamma \models \sigma$ 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 2.2.

Formally, $\Gamma \models \sigma$ iff

- (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 the induction hypothesis.)
- (2) (2) There are no v, τ, τ' such that $v \not\rightarrow y \in \Gamma$ and $\Gamma \models \tau := x \cdots \rightarrow v$ and $\Gamma \models \tau' := v \cdots \rightarrow u$. (τ may be the empty path, i.e. $x = v$.)
(σ itself is not precluded by split validity preclusion and a contradictory link. Note that $\tau \circ v \not\rightarrow y$ 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 valid)
or
there are $\Gamma \models \rho := x \cdots \rightarrow z$ and $\Gamma \models \rho' := z \cdots \rightarrow v$ for suitable ρ and ρ' .

- Case II.2: The negative case, i.e. σ a negative pp. $x \cdots \rightarrow u \not\rightarrow y$, $\sigma' := x \cdots \rightarrow u$, $\sigma = \sigma' \circ u \not\rightarrow y$ is entirely symmetrical.

Remark 2.1

(+++*** Orig. No.: Remark 2.1)

The following remarks all concern preclusion.

(1) Thus, in the case of preclusion, 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.

(2) Our definition is stricter than many usual ones, in the following sense: We require - according to our general picture to treat only direct links as information - that the preclusion "hits" the precluded path at the end, i.e. $v \not\rightarrow y \in \Gamma$, and ρ' hits $\tau \circ v \not\rightarrow y$ at v . In other definitions, it is possible that the preclusion hits at some v' , which is somewhere on the path τ , and not necessarily at its end. For instance, in the Tweety Diagram, see Diagram 2.1, if there were a node b' between b and d , we will need the path $c \rightarrow b \rightarrow b'$ to be valid, (obvious) validity of the arrow $c \rightarrow b$ will not suffice.

(3) If we allow ρ to be the empty path, then the case $z = x$ is a subcase of the present one.

(4) Our conceptual analysis has led to a very important simplification of the definition of validity. If we adopt on-path preclusion, we have to remember all paths which led to the information source to be considered: In the Tweety diagram, we have to remember that there is an arrow $a \rightarrow b$, it is not sufficient to note that we somehow came from a to b by a valid path, as the path $a \rightarrow c \rightarrow b \rightarrow d$ is precluded, but not the path $a \rightarrow b \rightarrow d$. If we adopt total path preclusion, we have to remember the valid path $a \rightarrow c \rightarrow b$ to see that it precludes $a \rightarrow c \rightarrow d$. If we allow preclusion to "hit" below the last node, we also have to remember the entire path which is precluded. Thus, in all those cases, whole paths (which can be very long) have to be remembered, but NOT in our definition.

We only need to remember (consider the Tweety diagram): (a) we want to know if $a \rightarrow b \rightarrow d$ is valid, so we have to remember a, b, d . Note that the (valid) path from a to b can be composed and very long. (b) we look at possible preclusions, so we have to remember $a \rightarrow c \not\rightarrow d$, again the (valid) path from a to c can be very long. (c) we have to remember that the path from c to b is valid (this was decided by induction before).

So in all cases (the last one is even simpler), we need only remember the starting node, a (or c), the last node of the valid paths, b (or c), and the information $\rightarrow d$ or $\not\rightarrow d$ - i.e. the size of what has to be recalled is ≤ 3 . (Of course, there may be many possible preclusions, but in all cases we have to look at a very limited situation, and not arbitrarily long paths.)

We take a fast look forward to Section 4.3, where we describe diagrams as information and its transfer, and nodes also as truth values. In these terms - and the reader is asked to excuse the digression - we may note above point (a) as $a \Rightarrow_b d$ - expressing that, seen from a , d holds with truth value b , (b) as $a \Rightarrow_c \neg d$, (c) as $c \Rightarrow_c b$ - and this is all we need to know.

□

We indicate here some modifications of the definition without discussion, which is to be found below.

(1) For on-path preclusion only: Modify condition (2) in Case II.1 to: (2') There is no v on the path σ (i.e. $\sigma : x \cdots \rightarrow v \cdots \rightarrow u$) such that $v \not\rightarrow y \in \Gamma$.

(2) For total validity preclusion: Modify condition (2) in Case II.1 to: (2') There are no v, τ, τ' such that $v \not\rightarrow y \in \Gamma$ and $\tau := x \cdots \rightarrow v$ and $\tau' := v \cdots \rightarrow u$ such that $\Gamma \models \tau \circ \tau'$.

(3) For extension based approaches: Modify condition (3) in Case II.1 as follows: (3') If there are conflicting paths, which are not precluded themselves by contradictory information, then we branch recursively (i.e. for all such situations) into two extensions, one, where the positive non-precluded paths are valid, one, where the negative non-precluded paths are valid.

Definition 2.5

(+++*** Orig. No.: Definition 2.5)

Finally, define $\Gamma \models xy$ iff there is $\sigma : x \rightarrow y$ s.th. $\Gamma \models \sigma$, likewise for $x\bar{y}$ and $\sigma : x \cdots \not\rightarrow y$.

Diagram 2.3 shows the most complicated situation for the positive case.

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

Fact 2.2

(+++*** Orig. No.: Fact 2.2)

The above definition and the one outlined in Definition 2.2 correspond.

Proof:

As Definition 2.2 is informal, this cannot be a formal proof, but it is obvious how to transform it into one.

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

Consider then case (P2.2) in Definition 2.2, and start from some x .

Case 1:

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




Diagram 2.3

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.

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 Definition 2.2.

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 2.2 and II.1 (2) + (3) correspond. Note that variant (P2.1) of Definition 2.2 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, we do not add any other information either in the approach in Definition 2.2.

Thus, both approaches are equivalent.

□

2.3 Review of other approaches and problems

We now discuss shortly in more detail some of the differences between various major definitions of inheritance formalisms.

Diagram 6.8, *p.* 179, in [Sch97-2] (which is probably due to folklore of the field) shows requiring downward chaining would be wrong. We repeat it here, see Diagram 2.4.

The problem of downward chaining:




Diagram 2.4

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*".) See the remark in Section 4.4.

Extension-based versus directly skeptical definitions

As this distinction has already received detailed discussion in the literature, we shall 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, the intersection of their conclusion sets). The seminal presentation is that of [Tou86], as refined by [San86]. 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 [HTT87]. 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 $x\bar{y}$, 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.

It has been shown in [Sch93] that the intersection of extensions is fundamentally different from the directly sceptical approach. See the remark in Section 4.4.

From now on, all definitions considered shall be (at least) upward chaining.

On-path versus off-path preclusion

This is a rather technical distinction, discussed in [THT87]. Briefly, a path $\sigma: x \rightarrow \dots \rightarrow y \rightarrow \dots \rightarrow z$ and a direct link $y \not\rightarrow u$ is an off-path preclusion of $\tau: x \rightarrow \dots \rightarrow z \rightarrow \dots \rightarrow u$, but an on-path preclusion only iff all nodes of τ between x and z lie on the path σ .

For instance, in the Tweety diagram, the arrow $c \not\rightarrow d$ is an on-path preclusion of the path $a \rightarrow c \rightarrow b \rightarrow d$, but the paths $a \rightarrow c$ and $c \rightarrow b$, together with $c \not\rightarrow d$, is an (split validity) off-path preclusion of the path $a \rightarrow b \rightarrow c$.

Split-validity versus total-validity preclusion

Consider again a preclusion $\sigma : u \rightarrow \dots \rightarrow x \rightarrow \dots \rightarrow v$, and $x \not\rightarrow y$ of $\tau : u \rightarrow \dots \rightarrow v \rightarrow \dots \rightarrow y$. 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 ([GV89], [KK89], [KKW89a], [KKW89b]) content themselves with the combinatorially simpler separate (split) validity of the lower and upper parts of σ : $\sigma' : u \rightarrow \dots \rightarrow x$ and $\sigma'' : x \rightarrow \dots \rightarrow v$. In Diagram 2.5, taken from [Sch97-2], 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 2.5

Thus, split validity preclusion will give here the definite result $u\bar{y}$. With total validity preclusion, the diagram has essentially the form of a Nixon Diamond.

3 Introduction to small and big sets and the logical systems P and R

It is natural to interpret "normality" by some sort of "size": "normality" might just mean "majority" (perhaps with different weight given to different cases), or something like "a

big subset". The standard abstraction of "big" is the notion of a filter (or, dually, an ideal is the abstraction of "small"). We include immediately a modification, the weak versions, to be discussed below. They seem to be minimal in the following sense: A reasonable abstract notion of size without the properties of weak filters seems difficult to imagine: The full set seems the best candidate for a "big" subset, "big" should cooperate with inclusion, and, finally, no set should be big and small at the same time.

Definition 3.1

(+++*** Orig. No.: Definition 3.1)

Fix a base set X .

A (weak) filter on or over X is a set $\mathcal{F} \subseteq \mathcal{P}(X) - \mathcal{P}(X)$ the power set of X -, such that (F1) – (F3) ((F1), (F2), (F3')) respectively) hold:

(F1) $X \in \mathcal{F}$

(F2) $A \subseteq B \subseteq X, A \in \mathcal{F}$ imply $B \in \mathcal{F}$

(F3) $A, B \in \mathcal{F}$ imply $A \cap B \in \mathcal{F}$

(F3') $A, B \in \mathcal{F}$ imply $A \cap B \neq \emptyset$.

So a weak filter satisfies (F3') instead of (F3).

An (weak) ideal on or over X is a set $\mathcal{I} \subseteq \mathcal{P}(X)$, such that (I1) – (I3) ((I1), (I2), (I3')) respectively) hold: (I1) $\emptyset \in \mathcal{I}$

(I2) $A \subseteq B \subseteq X, B \in \mathcal{I}$ imply $A \in \mathcal{I}$

(I3) $A, B \in \mathcal{I}$ imply $A \cup B \in \mathcal{I}$

(I3') $A, B \in \mathcal{I}$ imply $A \cup B \neq X$.

So a weak ideal satisfies (I3') instead of (I3).

Elements of a filter on X are called big subsets of X , their complements are called small, and the rest have "medium size". The set of the X -complements of the elements of a filter form an ideal, and vice versa. Note that these notions of "big" and "small" are by definition relative to the base set X . This is a different, additional notion of relativity than the one we will see in Section 5.

These notions are related to nonmonotonic logics as follows:

We can say that, normally, ϕ implies ψ iff in a big subset of all ϕ -cases, ψ holds. In preferential terms, ϕ implies ψ iff ψ holds in all minimal ϕ -models. If μ is the model choice function of a preferential structure, i.e. $\mu(\phi)$ is the set of minimal ϕ -models, then $\mu(\phi)$ will be a (the smallest) big subset of the set of ϕ -models, and the filter over the ϕ -models is the principal filter generated by $\mu(\phi)$.

Due to the finite intersection property, filters and ideals work well with logics: If ϕ holds normally, as it holds in a big subset, and so does ϕ' , then $\phi \wedge \phi'$ will normally hold, too, as the intersection of two big subsets is big again. This is a nice property, but not justified in all situations, consider e.g. simple counting of a finite subset. (The question has a name, "lottery paradox": normally no single participant wins, but someone wins in the end.) This motivates the weak versions.

Normality defined by (weak or not) filters is a local concept: the filter defined on X and the one defined on X' might be totally independent. Consider, however, the following two situations: Let Y' be a big subset of X' , $X \subseteq X'$, and $Y' \subseteq X$. If "size" has any absolute meaning, then Y' should be a big subset of X , too. On the other hand, let X and X' be big subsets of Y , then there are good reasons (analogue to those justifying the intersection property of filters) to assume that $X \cap X'$ is also a big subset of X' . These set properties are strongly connected to logical properties: For instance, if the latter property holds, we can deduce the logical property Cautious Monotony (see below for a formal definition): If ψ implies normally ϕ and ϕ' , because the sets X and X' of $\psi \wedge \phi$ -models and $\psi \wedge \phi'$ -models are big subsets of the set Y of ψ -models, then $\psi \wedge \phi'$ will imply normally ϕ too, as the set $X \cap X'$ of $\psi \wedge \phi \wedge \phi'$ -models will be a big subset of the set X' of $\psi \wedge \phi'$ -models.

More precisely, the reasoning is (a little simplified) as follows: If A and A' are small subsets of B , then A will also be a small subset of $B - A'$. Changing the reference set B just a little will not affect size. (In more detail: If A and A' are small subsets of B , then $A - A'$ will also be a small subset of $B - A'$.) Note that "small" is used here in two conceptually very different ways: (1) $\psi \sim \phi$ iff the set of $\psi \wedge \neg\phi$ -models is a small subset of the ψ -models, (2) when we change the reference class B just a little bit, relative size is preserved.

Seen more abstractly, such set properties allow the transfer of big subsets from one to another base set (and the conclusions drawn on this basis), and we call them "coherence properties". They are very important, not only for working with a logic which respects them, but also for soundness and completeness questions, often they are at the core of such problems. The reader is invited to read the articles by Ben-David and Ben-Eliyahu [BB94] and Friedman and Halpern [FH98], which treat essentially the same questions in different languages (and perhaps their comparison by the author in [Sch97-4] and [Sch04]).

These notions are tied to logical properties as follows:

The left hand column presents the single formula version, the center column the theory version (a theory is, for us, an arbitrary set of formulas), the right hand column the algebraic version, describing the choice function on the model set, e.g. $f(X) \subseteq X$ corresponds to the rule $\phi \vdash \psi$ implies $\phi \sim \psi$ in the formula version, and to $\overline{T} \subseteq \overline{\overline{T}}$ in the theory version. A short discussion of some of the properties follows the table.

(PR) is also called infinite conditionalization - we choose the name for its central role for preferential structures. Note that in the presence of $(\mu \subseteq)$, and if \mathcal{Y} is closed under finite intersections, (μPR) is equivalent to

$$(\mu PR') \quad f(X) \cap Y \subseteq f(X \cap Y).$$

The system of rules (AND), (OR), (LLE), (RW), (SC), (CP), (CM), (CUM) is also called system P (for preferential), adding (RM) gives the system R (for rationality or rankedness).

(LLE), (RW), (CCL) will all hold automatically, whenever we work with fixed model sets. (SC) corresponds to the choice of a subset. (CP) is somewhat delicate, as it presupposes that the chosen model set is non-empty. This might fail in the presence of ever better choices, without ideal ones. (PR) is an infinitary version of one half of the deduction theorem: Let T stand for ϕ , T' for ψ , and $\phi \wedge \psi \vdash \sigma$, so $\phi \vdash \psi \rightarrow \sigma$, but $(\psi \rightarrow \sigma) \wedge \psi \vdash \sigma$. (CUM) (whose most interesting half in our context is (CM)) may best be seen as normal use of lemmas: We have worked hard and found some lemmas. Now we can take a rest, and come back again with our new lemmas. Adding them to the axioms will neither add new theorems, nor prevent old ones to hold. (RM) is perhaps best understood by looking at big and small subsets. If the set of $\phi \wedge \psi$ -models is a big subset of the set of ϕ -models, and the set of $\phi \wedge \psi'$ -models is a not a small subset of the set of ϕ -models (i.e. big or of medium size), then the set of $\phi \wedge \psi \wedge \psi'$ -models is a big subset of the set of $\phi \wedge \psi'$ -models. Again, we have a double use of small/big here.

Definition 3.2

(AND) $\phi \vdash \psi, \phi \vdash \psi' \Rightarrow$ $\phi \vdash \psi \wedge \psi'$	(AND) $T \vdash \psi, T \vdash \psi' \Rightarrow$ $T \vdash \psi \wedge \psi'$	
(OR) $\phi \vdash \psi, \phi' \vdash \psi \Rightarrow$ $\phi \vee \phi' \vdash \psi$	(OR) $T \vdash \psi, T' \vdash \psi \Rightarrow$ $T \vee T' \vdash \psi$	$(\mu \cup w) - w$ for weak $f(A \cup B) \subseteq f(A) \cup f(B)$
(LLE) or Left Logical Equivalence $\vdash \phi \leftrightarrow \phi', \phi \vdash \psi \Rightarrow$ $\phi' \vdash \psi$	(LLE) $\overline{T} = \overline{T'} \Rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$	
(RW) or Right Weakening $\phi \vdash \psi, \vdash \psi \rightarrow \psi' \Rightarrow$ $\phi \vdash \psi'$	(RW) $T \vdash \psi, \vdash \psi \rightarrow \psi' \Rightarrow$ $T \vdash \psi'$	
(CCL) or Classical Closure	(CCL) $\overline{\overline{T}}$ is classically closed	
(SC) or Supraclassicality $\phi \vdash \psi \Rightarrow \phi \vdash \psi$	(SC) $\overline{T} \subseteq \overline{\overline{T}}$	$(\mu \subseteq)$ $f(X) \subseteq X$
(CP) or Consistency Preservation $\phi \vdash \perp \Rightarrow \phi \vdash \perp$	(CP) $T \vdash \perp \Rightarrow T \vdash \perp$	$(\mu \emptyset)$ $f(X) = \emptyset \Rightarrow X = \emptyset$
(RM) or Rational Monotony $\phi \vdash \psi, \phi \not\vdash \psi' \Rightarrow$ $\phi \wedge \psi' \vdash \psi$	(RM) $T \vdash \psi, T \not\vdash \psi' \Rightarrow$ $T \cup \{\psi'\} \vdash \psi$	$(\mu =)$ $X \subseteq Y, Y \cap f(X) \neq \emptyset \Rightarrow$ $f(X) = f(Y) \cap X$
(CM) or Cautious Monotony $\phi \vdash \psi, \phi \vdash \psi' \Rightarrow$ $\phi \wedge \psi \vdash \psi'$	(CM) $T \subseteq \overline{T'} \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T}} \subseteq \overline{\overline{T'}}$	$f(X) \subseteq Y \subseteq X \Rightarrow$ $f(Y) \subseteq f(X)$
(CUM) or Cumulativity $\phi \vdash \psi \Rightarrow$ $(\phi \vdash \psi' \Leftrightarrow \phi \wedge \psi \vdash \psi')$	(CUM) $T \subseteq \overline{T'} \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T}} = \overline{\overline{T'}}$	(μCUM) $f(X) \subseteq Y \subseteq X \Rightarrow$ $f(Y) = f(X)$
$\overline{\overline{\phi \wedge \phi'}} \subseteq \overline{\overline{\phi}} \cup \overline{\overline{\{\phi'\}}}$	(PR) $\overline{\overline{T \cup T'}} \subseteq \overline{\overline{\overline{T} \cup T'}}$	$(\mu \overline{PR})$ $X \subseteq Y \Rightarrow$ $f(Y) \cap X \subseteq f(X)$

We recall that smooth preferential models generate the system P , whereas smooth ranked models generate the system R - see e.g. [KLM90] or [Sch04] for details.

4 Interpretations

4.1 Introduction

We will discuss in this Section three interpretations of inheritance nets.

First, we will indicate fundamental differences between inheritance and the systems P and R , they will be elaborated in Section 5, where an interpretation in terms of small sets will be tried nonetheless.

Second, we will interpret inheritance nets as systems of information and information flow.

Third, we will interpret inheritance nets as systems of prototypes.

Inheritance nets present many intuitively attractive properties, thus it is not surprising that we can interpret them in several ways. Similarly, preferential structures can be used as a semantics for deontic, and for nonmonotonic logic, they express a common idea: choosing a subset of models by a binary relation. Thus, such an ambiguity need not be a sign for a basic flaw.

4.2 Informal comparison of inheritance with the systems P and R

The main issues

In the author's opinion, the following two properties of inheritance diagrams show the deepest difference to preferential and similar semantics, and the first even to classical logic. They have to be taken seriously, as they are at the core of inheritance systems, are independent of the particular formalism, and show that there is a fundamental difference between the former and the latter. Consequently, any attempt at translation will have to stretch one or both sides perhaps beyond the breaking point.

- (1) Relevance,
- (2) subideal situations, or relative normality

Both (and more) can be illustrated by the following simple Diagram 4.1 (which also shows conflict resolution by specificity).

(1) Relevance: As there is no monotonous path whatever between e and d , the question whether e 's are d 's or not, or vice versa, does not even arise. For the same reason, there is no question whether b 's are c 's, or not. (As a matter of fact, we will see below that b 's are non- c 's in system P - see Definition 3.2). In upward chaining formalisms, as there is no valid positive path from a to d , there is no question either whether a 's are f 's or not.

The problem of relevance has a trivial answer in upward chaining inheritance nets: there

must be some valid initial segment in the diagram to establish relevance.

Of course, in classical logic, all information is relevant to the rest, so we can say e.g. that e 's are d 's, or e 's are *non* - d 's, or some are d 's, some are not, but there is a connection. As preferential models are based on classical logic, the same argument applies to them.

(2) In our diagram, a 's are b 's, but not ideal b 's, as they are not d 's, the more specific information from c wins. But they are e 's, as ideal b 's are. So they are not perfectly ideal b 's, but as ideal b 's as possible. Thus, we have graded ideality, which does not exist in preferential and similar structures. In those structures, if an element is an ideal element, it has all properties of such, if one such property is lacking, it is not ideal, and we *can't* say anything any more. Here, however, we sacrifice as little normality as possible, it is thus a minimal change formalism.

In comparison, questions of information transfer and strength of information seem lesser differences. Already systems P and R (see Definition 3.2) differ on information transfer. In both cases, transfer is based on the same notion of smallness, which describes ideal situations. But, as said in Section 3, conceptually, this is very different from the use of "ideal-smallness", describing normal situations. Thus, it can be considered also on this level an independent question, and we can imagine systems based on absolutely ideal situations for normality, but with a totally different transfer mechanism.

For these reasons, extending preferential and related semantics to cover inheritance nets seems to stretch them to the breaking point, Thus, we should also look for other interpretations. (The term "interpretation" is used here in a non-technical sense.) In particular, it seems worth while to connect inheritance systems to other problems, and see whether there are similarities there. This is what we do now. We come back to the systems P and R in Section 5.

4.3 Inheritance as information transfer

An informal argument showing parallel ideas common to inheritance with an upward chaining formalism and information transfer is as follows: First, arrows represent certainly some kind of information, of the kind "most a 's are b 's" or so. (See Diagram 4.1.) Second, to be able to use information, e.g. " d 's are f 's" at a , we have to be able to connect from a to d by a valid path, this information has to be made accessible to a , or, in other terms, a working information channel from a to d has to be established. Third, specificity (when present) decides conflicts (we take the split validity approach). This can be done procedurally, or, perhaps simpler and certainly in a more transparent way, by assigning a comparison of information strength to valid paths. Now, information strength may also be called truth value (to use a term familiar in logic) and the natural entity at hand is the node itself - this is just a cheap formal trick without any conceptual meaning.

When we adopt this view, nodes and arrows (and valid paths) have multiple functions,




Diagram 4.1

and it may seem that we overload the (deceptively) simple picture. But it is perhaps the charm and the utility and naturalness of inheritance systems that they are not "clean", and hide many complications under a simple surface, as human common sense reasoning often does, too.

In a certain way, this is a poor man's interpretation, as it does not base inheritance on another formalism, but gives only an intuitive reading. Yet, it gives a connection to other branches of reasoning, and is as such already justified - in the author's opinion. Moreover, our analysis makes a clear distinction between arrows and composite valid paths. This distinction is implicit in inheritance formalisms, we make it explicit through our concepts.

But this interpretation is by no means the only one, and can only be suggested as a possibility.

We will now first give the details, and then discuss our interpretation.

(1) Information:

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" - and $A \not\rightarrow B$ will mean that most A 's are not B 's.

(2) Information sources and flow:

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".

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

(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) Information strength:

A valid, composed or atomic positive path σ from A' to A allows us 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' . 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. Of course, if desired, we can also adopt this (much more complicated) idea.

Nodes are also truth values. They are the strength of the information whose source they are. (This might seem an abuse of nodes, but we already have them, so why not use them?)

Discussion:

Considering direct arrows as information meets probably with little objection.

The conclusion of a valid path (e.g. if $\sigma : a \rightarrow b$ is valid, then its conclusion is " a 's are b 's") is certainly also information, but it has a status different from the information of a direct link, so we should distinguish it clearly. At least in upward chaining formalisms, using the path itself as some channel through which information flows, and not the conclusion, seems more natural. The conclusion says little about the inner structure of the path, which is very important in inheritance networks, e.g. for preclusion. When calculating validity of paths, we look at (sub- and other) paths, but not their results, and should also

express this clearly.

Once we accept this picture of valid positive paths as information channels, it is natural to see their upper ends as information sources.

Our interpretation supports upward chaining, and vice versa, upward chaining supports our interpretation.

One of the central ideas of inheritance is preclusion, which, in the case of split validity preclusion, works by an absolute comparison between nodes. Thus, if we accept split validity preclusion, it is natural to see valid positive paths as comparisons between information of different strengths. Conversely, if we accept absolute comparison of information, we should also accept split validity preclusion - these interpretations support each other.

Whatever type of preclusion we accept, preclusion clearly compares information strength, and allows us to decide for the stronger one. We can see this procedurally, or by giving different values to different information (depending on the source), which we can call truth values to connect our picture to other areas of logic. It is then natural - as we have it already - to use the source node itself as truth value, with comparison via valid positive paths.

Illustration:

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 Tweety diagram, see Diagram 2.1, 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.

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. A valid negative composed path has no function, only its parts have.

We obtain automatically 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.

Our interpretation underlines the importance of initial segments: Initial segments make information accessible. Thus, initial segments have to be valid.

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.

Further comments:

Our reading also covers enriched diagrams, where arbitrary information can be "appended" to a node.

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 something, 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 believes and U' does NOT believe some information of A , then either U' 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.)

The restriction that negative links can only be information 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, perhaps without any real gain in insight.

We translate the analysis and decision of Definition 2.2 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 with respect to B , then the weaker information is discarded.
- (3) There remains a (nonempty, by lack of cycles) set of the A_i , such that for no such A_i there is A_j with better 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.

We continue now Remark 2.1, (4), and turn this into a formal system.

Fix a diagram Γ , and do an induction as in Definition 2.2.

Definition 4.1

(+++*** Orig. No.: Definition 4.1)

- (1) We distinguish $a \Rightarrow b$ and $a \Rightarrow_x b$, where the intuition of $a \Rightarrow_x b$ is: we know with strength x that a's are b's, and of $a \Rightarrow b$ that it has been decided taking all information into consideration that $a \Rightarrow b$ holds.
- (2) $a \rightarrow b$ implies $a \rightarrow_a b$, likewise $a \not\rightarrow b$ implies $a \rightarrow_a \neg b$.
- (3) $a \Rightarrow_a b$ implies $a \Rightarrow b$, likewise $a \Rightarrow_a \neg b$ implies $a \Rightarrow \neg b$. This expresses the fact that direct arrows are uncontested.
- (4) $a \Rightarrow b$ and $b \Rightarrow_b c$ imply $a \Rightarrow_b c$, likewise for $b \Rightarrow_b \neg c$. This expresses concatenation - but without deciding if it is accepted! Note we cannot make $a \Rightarrow b$ and $b \Rightarrow c$ imply $a \Rightarrow_b c$ a rule, as this would make concatenation of two composed paths possible.
- (5) We decide acceptance of composed paths as in Definition 2.2, where preclusion uses accepted paths for deciding.

Note that we also reason in this system about relative strength of truth values, which are just nodes, this is then, of course, used in the acceptance condition, in preclusion, more precisely.

4.4 Inheritance as reasoning with prototypes

Some of the issues we discuss here apply also to the more general picture of information and its transfer. We present them here for motivational reasons: it seems easier to discuss them in the (somewhat!) more concrete setting of prototypes than in the very general situation of information handling. These issues will be indicated.

It seems natural to see information in inheritance networks as information about prototypes. (We do not claim that our use of the word "prototype" has more than a vague relation to the use in psychology. We do not try to explain the usefulness of prototypes either, one possibility is that there are reasons why birds fly, and why penguins *don't*, etc.) In the Tweety diagram, we will thus say that prototypical birds will fly, prototypical penguins will not fly. More precisely, the property "fly" is part of the bird prototype, the property " \neg fly" part of the penguin prototype. Thus, the information is given for some node, which defines its application or domain (bird or penguin in our example) - beyond this node, the property is not defined (unless inherited, of course). 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. This again shows that we are very far from the ideal picture of small and big subsets as used in systems P and R . (This, of course, goes beyond the problem of prototypes.)

Of course, we will want to "inherit" properties of prototypes, for instance in Diagram 4.1, a "should" inherit the property e from b , and the property $\neg d$ from c . Informally, we will argue as follows: "Prototypical a's have property b , and prototypical b's have property e ,

so it seems reasonable to assume that prototypical a 's also have property e - unless there is better information to the contrary." A plausible solution is then to use upward chaining inheritance as described above to find all relevant information, and then compose the prototype.

We discuss now three points whose importance goes beyond the treatment of prototypes:

(1) Using upward chaining has an additional intuitive appeal: We consider information at a the best, so we begin with b (and c), and only then, tentatively, add information e from b . Thus, we begin with strongest information, and add weaker information successively - this seems good reasoning policy.

(2) In upward chaining, we also collect information at the source (the end of the path), and do not use information which was already filtered by going down - thus the information we collect has no history, and we cannot encounter problems of iterated revision, which are problems of history of change. (In downward chaining, we only store the reasons why something holds, but not why something does not hold, so we cannot erase this negative information when the reason is not valid any more. This is an asymmetry apparently not much noted before. Consider Diagram 2.4. Here, the reason why u does not accept y as information, but $\neg y$, is the preclusion via x . But from z , this preclusion is not valid any more, so the reason why y was rejected is not valid any more, and y can now be accepted.)

(3) We come back to the question of extensions vs. direct scepticism. Consider the Nixon Diamond, Diagram 2.2. Suppose Nixon were a subclass of Republican and Quaker. Then the extensions approach reasons as follows: Either the Nixon class prototype has the pacifist property, or the hawk property, and we consider these two possibilities. But this is not sufficient: The Nixon class prototype might have neither property - they are normally neither pacifists, nor hawks, but some are this, some are that. So the conceptual basis for the extensions approach does not hold: "Tertium non datur" just simply does not hold (as in Intuitionist Logic, where we may have neither a proof for ϕ , nor for $\neg\phi$).

Once we fixed this decision, i.e. how to find the relevant information, we can still look upward or downward in the net and investigate the changes between the prototypes in going upward or downward, as follows: E.g., in above example, we can look at the node a and its prototype, and then at the change going from a to b , or, conversely, look at b and its prototype, and then at the change going from b to a . The problem of finding the information, and this dynamics of information change have to be clearly separated.

In both cases, we see the following:

(1) The language is kept small, and thus efficient.

For instance, when we go from a to b , information about c is lost, and " c " does not figure any more in the language, but f is added. When we go from b to a , f is lost, and c is won. In our simple picture, information is independent, and contradictions are always between two bits of information.

(2) Changes are kept small, and need a reason to be effective. Contradictory, stronger information will override the old one, but no other information, except in the following case: making new information (in-) accessible will cause indirect changes, i.e. information now made (in-) accessible via the new node. This is similar to formalisms of causation: if a reason is not there any more, its effects vanish, too.

It is perhaps more natural when going downward also to consider "subsets", as follows: Consider Diagram 4.1. b 's are d 's, and c 's are $\neg d$'s, and c 's are also b 's. So it seems plausible to go beyond the language of inheritance nets, and conclude that b 's which are not c 's will be d 's, in short to consider $(b - c)$'s. It is obvious which such subsets to consider, and how to handle them: For instance, loosely speaking, in $b \cap d$ e will hold, in $b \cap c \cap d$ $\neg f$ will hold, in $b \cap d \cap \mathbf{C}c$ f will hold, etc. This is just putting the bits of information together.

We turn to another consideration, which will also transcend the prototype situation and we will (partly) use the intuition that nodes stand for sets, and arrows for (soft) inclusion in a set or its complement.

In this reading, specificity stands for soft set inclusion. So, if b and c are visible from a , and there is a valid path from c to b (as in Diagram 4.1), then a is a subset both of b and c , and c a subset of b , so $a \subseteq c \subseteq b$ (softly). But then a is closer to c than a is to b . Automatically, a will be closest to itself.

When we go now from b to c , we loose information d and f , win information $\neg d$, but keep information e . Thus, this is minimal change: we give up (and win) only the necessary information, but keep the rest. As our language is very simple, we can use the Hamming distance between formula sets here. (We will make a remark on more general situations just below.)

When we look now again from a , we take the set-closest class (c), and use the information of c , which was won by minimal change (i.e. the Hamming closest) from information of b . So we have the interplay of two distances, where the set distance certainly is not symmetrical, as we need valid paths for access and comparison. If there is no such valid path, it is reasonable to make the distance infinite.

The promised remark on more general situations: in richer languages, we cannot count formulas to determine the Hamming distance between two situations (i.e. models or model sets), but have to take the difference in propositional variables. Consider e.g. the language with two variables, p and q . The models (described by) $p \wedge q$ and $p \wedge \neg q$ have distance 1, whereas $p \wedge q$ and $\neg p \wedge \neg q$ have distance 2. Note that this distance is NOT robust under re-definition of the language. Let p' stand for $(p \wedge q) \vee (\neg p \wedge \neg q)$, and q' for q . Of course, p' and q' are equivalent descriptions of the same model set, as we can define all the old singletons also in the new language. Then the situations $p \wedge q$ and $\neg p \wedge \neg q$ have now distance 1, as one corresponds to $p' \wedge q'$, the other to $p' \wedge \neg q'$.

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.)

We conclude this Section with two remarks, the first on prototypes, the second a general remark on preclusion in inheritance systems.

Realistic prototypical reasoning will probably neither always be upward, nor always be downward. A medical doctor will not begin with the patient's most specific category (name and birthday or so), nor will he begin with all he knows about general objects. Therefore, it seems reasonable to investigate upward and downward reasoning here.

Obviously, in some 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.

5 Detailed translation of inheritance to modified systems of small sets

5.1 Normality

As we saw already in Section 4.2, normality in inheritance (and Reiter defaults etc.) is relative, and as much normality as possible is preserved. There is no absolute $N(X)$, but only $N(X, \phi)$, and $N(X, \psi)$ might be defined, but not $N(X, \psi)$. Normality in the sense of preferential structures is absolute: if x is not in $N(X)$ ($= \mu(X)$ in preferential reading), 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 with respect to 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. 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.

What are the laws of relative normality? $N(X, \phi)$ and $N(X, \psi)$ will be largely independent (except for trivial situations, where $\phi \leftrightarrow \psi$, ϕ is a tautology, etc.). $N(X, \phi)$ 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, as illustrated by Fact 5.1.

We interpret now "normal" by "big set", i.e. essentially " ϕ holds normally in X iff there is a big subset of X , where ϕ holds". This will, of course, be modified.

5.2 Small sets

The main interest of this Section is perhaps to show the adaptations of the concept of small/big subsets necessary for a more "real life" situation, where we have to relativize. The amount of changes we do illustrates the problems and what can be done, but also perhaps what should not be done, as the concept is stretched too far.

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.)

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 essentially independent, may perhaps even have an empty intersection, only one may be defined, etc.

Consider now the system P (with Cumulativity), see Definition 3.2. Small sets (see Definition 3.1) 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 Cumulativity, or Cautious Monotony more precisely: if the set of $\alpha \wedge \neg\gamma$ -cases is a small subset of the set of α -cases, then $\alpha \sim \beta$ carries over to $\alpha \wedge \gamma : \alpha \wedge \gamma \sim \beta$. (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.

Fact 5.1

(+++*** Orig. No.: Fact 5.1)

- (a) System P is too strong to capture inheritance.
- (b) System R is too weak to capture inheritance.

Proof:

(a) Consider the Tweety diagram, Diagram 2.1. $c \rightarrow b \rightarrow d, c \not\rightarrow d$. There is no arrow $b \not\rightarrow c$, 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 $c \cap b$ is a small subset of b , which we write $c \cap b < b$. $c \cap b = (c \cap b \cap d) \cup (c \cap b \cap \mathbf{C}d)$. $c \cap b \cap \mathbf{C}d \subseteq b \cap \mathbf{C}d < b$, the latter by $b \rightarrow d$, thus $c \cap b \cap \mathbf{C}d < b$, essentially by Right Weakening. Set now $X := c \cap b \cap d$. As $c \not\rightarrow d$, $X := c \cap b \cap d \subseteq c \cap d < c$, and by the same reasoning as above $X < c$. It remains to show $X < b$. We use now $c \rightarrow b$. As $c \cap \mathbf{C}b < c$, and $c \cap X < c$, by Cumulativity $X = c \cap X \cap b < c \cap b$, so essentially by OR $X = c \cap X \cap b < b$. Using the filter property, we see that $c \cap b < b$.

(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 .

□

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. This corresponds to the distinction between $a \Rightarrow b$ and $a \Rightarrow_a b$ in Definition 4.1.

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].

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 to preclusion, as it should be - see below.

For the situation $X \rightarrow Y \rightarrow Z$, we will then 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 . (We generalize already to the case where there is a valid path from X to Y .)

We call this procedure information transfer.

$Y \rightarrow Z$ 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.

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. The future will tell if any of the ideas can be used in other contexts.

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.

(In slight generalization:) 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 right), 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 Cumulativity or even Rational Monotony, 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 with respect to 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 with respect to 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.

For (2):

By $X \rightarrow Y$, X and Y are related, and we assume that X behaves as $Y \cap X$ does with respect to 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 with respect to 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 .

(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 Diagram 5.1:

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




Diagram 5.1

U is not considered), and which is accessible (thus Y'' is not considered). (We can be slightly 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 , then 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 $Y \cap Y''$.

$Y' \cap CZ$ will be a Z -BIG subset of Y' , $Y \cap Y' \cap CZ$ 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. This is the argument for the simplified definition of preclusion - see Remark 2.1, (4).)

(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 $Y \rightarrow Z$), $Y' \cap CZ$ is a Z -BIG subset of Y' (by $Y' \rightarrow Z$), so by preclusion $Y \cap Y' \cap CZ$ is a Z -BIG subset of $Y \cap Y'$.

(1.4) $Y \cap Y' \cap CZ$ 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 $X \cap Y \cap Y' \cap CZ$ 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 $X \cap Y \cap Y' \cap CZ$ 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 $X \cap CZ$ 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 $Y' \cap CZ \in B(Y', Z) \Rightarrow Y \cap Y' \cap CZ \in B(Y \cap Y', Z)$ by preclusion

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

(2') $X \cap Y \cap Y' \cap CZ \in b(X \cap Y \cap Y', Z)$ and $X \cap Y \cap Y' \in b(X, Y \cap Y') \Rightarrow X \cap CZ \in b(X, Z)$.

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 conclusion:

(2.1) All contradictory paths are out: e.g. $Y \cap Z$ will be Z-big, but $Y \cap Y' \cap CZ$ 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 Diagram 5.2:




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, $Y' \cap CZ$ is Z-BIG, $Y'' \cap Z$ is Z-BIG, and $Y \cap Y' \cap CZ$ 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 5.2

(+++*** Orig. No.: Fact 5.2)

The above definition and the one outlined in Definition 2.3 correspond.

Proof:

By induction on the length of the deduction that $X \cap Z$ (or $X \cap CZ$) 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 $X \cap CY \in B(X, Y)$ 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 $X \cap CZ$.

We stay close now to above Diagram 5.1, so we argue for the negative case.

Suppose that we have deduced $X \cap CZ \in b(X, Z)$, 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 $X \cap \cap Y_i \in b(X, \cap Y_i)$, 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 $X \cap \cap Y_i \in b(X, \cap Y_i)$ 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 the induction hypothesis, there are valid paths from X to the Y_i .

The first hypothesis of (2'), $X \cap \cap Y_i \cap CZ \in b(X \cap \cap Y_i, Z)$ 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 $\cap Y_i \cap CZ \in B(\cap Y_i, Z)$ 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 with respect to X and $Z!$), and we are done by the definition of valid paths in Section 2.

□

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 the systems P and R are, as we have to transfer information also to perhaps small subsets. We had to be more careful, as P and 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.

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 of combining different small/big sets. This is, of course, a special form of relevance, there might be other forms.

6 ACKNOWLEDGEMENTS

(+++*** Orig.: ACKNOWLEDGEMENTS) Some of above reflections originated directly or indirectly from a long email discussion and criticisms with and from A.Bochman. Indirectly, the interest of my former PhD student J.Ben-Naim in multi-valued and information source logic has certainly influenced me to reconsider inheritance networks from these perspectives. Three referees have helped with their very constructive criticism to make this text readable.

References

- [ABK07] A.Avron, J.Ben-Naim, B.Konikowska: "Cut-free ordinary sequent calculi for logics having generalized finite-valued semantics", Journal Logica Universalis, to appear
- [BB94] Shai Ben-David, R.Ben-Eliyahu: "A modal logic for subjective default reasoning", Proceedings LICS-94, 1994
- [BGH95] J.Barwise, D.Gabbay, C.Hartonas: "On the Logic of Information Flow", Journal of the IGPL, Vol. 3.1, 1995
- [FH98] N.Friedman, J.Halpern: "Plausibility measures and default reasoning", IBM Almaden Research Center Tech.Rept. 1995, to appear in Journal of the ACM
- [GV89] H.Geffner, T.Verma, "Inheritance = *Chaining* + *Defeat*", UCLA Technical Report CSD-890039, R-129-L, June 1989, also in: Proceed. Fourth International Symposium on Methodologies for Intelligent Systems, North Holland, 1989, p.411-418
- [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
- [KK89] T.Krishnaprasad, M.Kifer, "An Evidence-based Framework for a Theory of Inheritance", Proceed. IJCAI 89, p.1093-1098
- [KKW89a] T.Krishnaprasad, M.Kifer, D.Warren, "On the Circumscriptive Semantics of Inheritance Networks", Proceed. Fourth International Symposium on Methodologies for Intelligent Systems, North Holland, 1989, p.448-456
- [KKW89b] T.Krishnaprasad, M.Kifer, D.Warren, "On the Declarative Semantics of Inheritance Networks", Proceed. IJCAI 89, p.1098-1103
- [KLM90] S.Kraus, D.Lehmann, M.Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics", Artificial Intelligence, 44 (1-2), p.167-207, July 1990
- [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
- [MS91] D.Makinson, K.Schlechta, "Floating Conclusions and Zombie Paths", Artificial Intelligence (Research Note) 48 (1991) p. 199-209

- [Mor98] L.Morgenstern, "Inheritance comes of age: applying nonmonotonic techniques to problems in industry", *Artificial Intelligence*, vol. 103, pp. 1-34, 1998
- [SL89] B.Selman, H.Levesque, "The Tractability of Path-Based Inheritance", *Proceed. IJCAI 1989*, p.1140-1145
- [San86] E.Sandewall, "Non-monotonic inference rules for multiple inheritance with exceptions", *Proceedings IEEE 74* (1986), p.1345-1353
- [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
- [Sch97-4] K.Schlechta: "Filters and partial orders", *Journal of the Interest Group in Pure and Applied Logics*, Vol. 5, No. 5, p. 753-772, 1997
- [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