



HAL
open science

Defeasible inheritance systems and reactive diagrams

Dov Gabbay, Karl Schlechta

► **To cite this version:**

Dov Gabbay, Karl Schlechta. Defeasible inheritance systems and reactive diagrams. 2008. hal-00336105

HAL Id: hal-00336105

<https://hal.science/hal-00336105>

Preprint submitted on 31 Oct 2008

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.

Defeasible inheritance systems and reactive diagrams *

Dov M Gabbay †
King's College, London ‡

Karl Schlechta §
Laboratoire d'Informatique Fondamentale de Marseille ¶

November 1, 2008

Abstract

Inheritance diagrams are directed acyclic graphs with two types of connections between nodes: $x \rightarrow y$ (read x is a y) and $x \not\rightarrow y$ (read as x is not a y). Given a diagram D , one can ask the formal question of “is there a valid (winning) path between node x and node y ?” Depending on the existence of a valid path we can answer the question “ x is a y ” or “ x is not a y ”.

The answer to the above question is determined through a complex inductive algorithm on paths between arbitrary pairs of points in the graph.

This paper aims to simplify and interpret such diagrams and their algorithms. We approach the area on two fronts.

(1)

Suggest reactive arrows to simplify the algorithms for the winning paths.

(2)

We give a conceptual analysis of (defeasible or nonmonotonic) inheritance diagrams, and compare our analysis to the “small” and “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 of a particularly simple type.

We show that the small and 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.

We will also interpret inheritance diagrams as theories of prototypical reasoning, based on two distances: set difference, and information difference.

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

It is easily seen that inheritance diagrams can also be analysed in terms of reactive diagrams - as can all argumentation systems.

AMS Classification: 68T27, 68T30

Contents

1	Introduction	2
1.1	Terminology	2
1.2	Inheritance and reactive diagrams	2
1.3	Conceptual analysis	8

*Paper 326

†Dov.Gabbay@kcl.ac.uk, www.dcs.kcl.ac.uk/staff/dg

‡Department of Computer Science, King's College London, Strand, London WC2R 2LS, UK

§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	9
2.1	Basic discussion	9
2.2	Directly sceptical split validity upward chaining off-path inheritance	13
2.3	Review of other approaches and problems	19
3	Defeasible inheritance and reactive diagrams	21
3.1	Summary of our algorithm	21
3.2	Overview	22
3.3	Compilation and memorization	22
3.4	Executing the algorithm	23
3.5	Signposts	23
3.6	Beyond inheritance	25
4	Introduction to small and big sets and the logical systems P and R	25
5	Interpretations	34
5.1	Introduction	35
5.2	Informal comparison of inheritance with the systems P and R	35
5.3	Inheritance as information transfer	36
5.4	Inheritance as reasoning with prototypes	39
6	Detailed translation of inheritance to modified systems of small sets	41
6.1	Normality	41
6.2	Small sets	41
7	Conclusion	47
	References	48

karl-search= Start Is-In

LABEL: Section Is-In

1 Introduction

1.1 Terminology

“Inheritance” will stand here for “nonmonotonic or defeasible inheritance”. We will use indiscriminately “inheritance system”, “inheritance diagram”, “inheritance network”, “inheritance net”.

In this introduction, we first give the connection to reactive diagrams, then give the motivation, then describe in very brief terms some problems of inheritance diagrams, and mention the basic ideas of our analysis.

1.2 Inheritance and reactive diagrams

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 authors’ 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 the 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]).

An inheritance net is a directed graph with two types of connections between nodes $x \rightarrow y$ and $x \not\rightarrow y$. Diagram 2.2 (page 11) is such an example. The meaning of $x \rightarrow y$ is that x is also a y and the meaning of $x \not\rightarrow y$ is that x is not a y . We do not allow the combinations $x \not\rightarrow y \not\rightarrow z$ or $x \not\rightarrow y \rightarrow z$ but we do allow $x \rightarrow y \rightarrow z$ and $x \rightarrow y \not\rightarrow z$.

Given a complex diagram such as Diagram 2.3 (page 19) and two points say z and y , the question we ask is to determine from the diagram whether the diagram says that

- (1) z is y
- (2) z is not y
- (3) nothing to say.

Since in Diagram 2.3 (page 19) there are paths to y from z either through x or through v , we need to have an algorithm to decide. Let \mathcal{A} be such an algorithm.

We need \mathcal{A} to decide

- (1) are there valid paths from z to y
- (2) of the opposing paths (one which supports ‘ z is y ’ and one which supports ‘ z is not y ’), which one wins (usually winning makes use of being more specific but there are other possible options).

So for example, in Diagram 2.3 (page 19), the connection $x \rightarrow v$ makes paths though x more specific than paths through v . The question is whether we have a valid path from z to x .

In the literature, as well as in this paper, there are algorithms for deciding the valid paths and the relative specificity of paths. These are complex inductive algorithms, which may need the help of a computer for the case of the more complex diagrams.

It seems that for inheritance networks we cannot adopt a simple minded approach and just try to ‘walk’ on the graph from z to y , and depending on what happens during this ‘walk’ decide whether z is y or not.

To explain what we mean, suppose we give the network a different meaning, that of fluid flow. $x \rightarrow y$ means there is an open pipe from x to y and $x \not\rightarrow y$ means there is a blocked pipe from x to y .

To the question ‘can fluid flow from z to y in Diagram 2.3’ (page 19), there is a simple answer:

Fluid can flow iff there is a path comprising of \rightarrow only (without any $\not\rightarrow$ among them).

Similarly we can ask in the inheritance network something like (*) below:

- (*) z is (resp. is not) y according to diagram D , iff there is a path π from z to y in D such that some non-inductive condition $\psi(\pi)$ holds for the path π .

Can we offer the reader such a ψ ?

If we do want to help the user to ‘walk’ the graph and get an answer, we can proceed as one of the following options:

- Option 1. Add additional annotations to paths to obtain D^* from D , so that a predicate ψ can be defined on D^* using these annotations. Of course these annotations will be computed using the inductive algorithm in \mathcal{A} , i.e. we modify \mathcal{A} to \mathcal{A}^* which also executes the annotations.
- Option 2. Find a transformation τ on diagrams D to transform D to $D' = \tau(D)$, such that a predicate ψ can be found for D' . So we work on D' instead of on D .

We require a compatibility condition on options 1 and 2:


(C1) If we apply \mathcal{A}^* to D^* we get D^* again.

(C2) $\tau(\tau(D)) = \tau(D)$.

We now present the tools we use for our annotations and transformation. These are the reactive double arrows. Consider the following Diagram 1.1 (page 4):

Diagram 1.1 LABEL: Diagram Dov-Is-1

Reactive graph




We want to walk from a to e . If we go to c , a double arrow from the arc $a \rightarrow c$ blocks the way from d to e . So the only way to go to e is through b . If we start at a' there is no such block. It is the travelling through the arc (a, c) that triggers the double arrow $(a, c) \rightarrow (d, e)$.

We want to use \rightarrow in \mathcal{A}^* and in τ .

So in Diagram 2.3 (page 19) the path $u \rightarrow x \not\rightarrow y$ is winning over the path $u \rightarrow v \rightarrow y$, because of the specificity arrow $x \rightarrow v$. However, if we start at z then the path $z \rightarrow u \rightarrow v \rightarrow y$ is valid because of $z \not\rightarrow x$. We can thus add the following double arrows to the diagram. We get Diagram 1.2 (page 4).

Diagram 1.2 LABEL: Diagram Dov-Is-2



If we start from u , and go to $u \rightarrow v$, then $v \rightarrow y$ is cancelled. Similarly $u \rightarrow x \rightarrow v$ cancelled $v \rightarrow y$. So the only path is $u \rightarrow x \not\rightarrow y$.


If we start from z then $u \rightarrow x$ is cancelled and so is the cancellation $(u, v) \rightarrow (v, y)$. hence the path $z \rightarrow u \rightarrow v \rightarrow y$ is open.

We are not saying that $\tau(\text{Diagram1.1}) = \text{Diagram1.2}$ but something effectively similar will be done by τ .

We emphasize that the construction depends on the point of departure. Consider Diagram 2.3 (page 19) . Starting at u , we will have to block the path uvy . Starting at z , the path $zuvy$ has to be free. See Diagram 1.3 (page 5) .

Diagram 1.3 LABEL: Diagram U-D-reactive

The problem of downward chaining - reactive



So we cannot just add a double arrow from $u \rightarrow v$ to $v \rightarrow y$, blocking $v \rightarrow y$, and leave it there when we start from z . We will have to erase it when we change the origin.

At the same time, this shows an advantage over just erasing the arrow $v \rightarrow y$:

When we change the starting point, we can erase simply all double arrows, and do not have to remember the original diagram.

How do we construct the double arrows given some origin x ?

First, if all possible paths are also valid, there is nothing to do. (At the same time, this shows that applying the procedure twice will not result in anything new.)

Second, remember that we have an upward chaining formalism. So if a potential path fails to be valid, it will do so at the end.

Third, suppose that we have two valid paths $\sigma : x \rightarrow y$ and $\tau : x \rightarrow y$.

If they are negative (and they are either both negative or positive, of course), then they cannot be continued. So if there is an arrow $y \rightarrow z$ or $y \not\rightarrow z$, we will block it by a double arrow from the first arrow of σ and from the first arrow of τ to $y \rightarrow z$ ($y \not\rightarrow z$ respectively).

If they are positive, and there is an arrow $y \rightarrow z$ or $y \not\rightarrow z$, both $\sigma y \rightarrow z$ and $\tau y \rightarrow z$ are potential paths (the case $y \not\rightarrow z$ is analogue). One is valid iff the other one is, as σ and τ have the same endpoint, so preclusion, if present, acts on both the same way. If they are not valid, we block $y \rightarrow z$ by a double arrow from the first arrow of σ and from the first arrow of τ to $y \rightarrow z$.

Of course, if there is only one such σ , we do the same, there is just to consider to see the justification.


We summarize:

Our algorithm switches impossible continuations off by making them invisible, there is just no more arrow to concatenate. As validity in inheritance networks is not forward looking - validity of $\sigma : x \rightarrow y$ does not depend on what is beyond y - validity in the old and in the new network starting at x are the same. As we left only valid paths, applying the algorithm twice will not give anything new.

We illustrate this by considering Diagram 3.1 (page 24) .


First, we add double arrows for starting point c , see Diagram 1.4 (page 7) ,

Diagram 1.4 LABEL: Diagram I-U-reac-c



and then for starting point x , see Diagram 1.5 (page 7) .

Diagram 1.5 LABEL: Diagram I-U-reac-x



For more information on reactive diagrams, see also [Gab08b], and an earlier version: [Gab04].

1.3 Conceptual analysis

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, see the discussion in Section 2.3 (page 19). 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.

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].

The text is organized as follows. After an introduction to inheritance theory, connections with reactive diagrams in Section 3 (page 21) , and big and small subsets and the systems P and R in Section 2 (page 9) , we turn to an informal description of the fundamental differences between inheritance and the systems P and R in Section 5.2 (page 35) , give an analysis of inheritance systems in terms of information and information flow in Section 5.3 (page 36) , then in terms of reasoning with prototypes in Section 5.4 (page 39) , and conclude in Section 6 (page 41) with a translation of inheritance into (necessarily deeply modified) coherent systems of big and small sets, respectively logical systems P and R . One of the main modifications will be to relativize the notions of small and big, which thus become less “mathematically pure” but perhaps closer to actual use in “dirty” common sense reasoning.

karl-search= End Is-In

2 Introduction to nonmonotonic inheritance

karl-search= Start Is-Inh-Intro

LABEL: Section Is-Inh-Intro

2.1 Basic discussion

LABEL: Section IS-2.1

We give here an informal discussion. The reader unfamiliar with inheritance systems should consult in parallel Definition 2.3 (page 15) and Definition 2.4 (page 16) . As there are many variants of the definitions, it seems reasonable to discuss them before a formal introduction, which, otherwise, would seem to pretend to be definite without being so.

(Defeasible or nonmonotonic) inheritance networks or diagrams

(+++*** Orig.: (Defeasible or nonmonotonic) inheritance networks or diagrams)

LABEL: Section (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 IS-2.1)

LABEL: Definition IS-2.1

LABEL: Definition Inheritance-Net

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” for “normal”, thus “most elements of X are in Y ”,

"most elements of X are not in Y ", etc. These are by no means the only interpretations, as we will see - we will use these expressions for the moment just to help the reader's intuition. We should add immediately a word of warning: "most" is here not necessarily, but only by default, transitive, in the following sense. In the Tweety diagram, see Diagram 2.1 (page 10) below, most penguins are birds, most birds fly, but it is not the case that most penguins fly. This is the problem of transfer of relative size which will be discussed extensively, especially in Section 6 (page 41) .


According to the set interpretation, we will also use informally expressions like $X \cap Y$, $X - Y$, CX - where C stands for set complement -, 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 here 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 $penguins \not\rightarrow fly$ 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 (page 10) , 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.)

Diagram 2.1 LABEL: *Diagram Tweety*

The Tweety diagram



(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 (page 10) . 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 (nonempty) 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 (page 10) . Then, the concept of preclusion collapses, as there are now equivalent arguments to accept $a \rightarrow b \rightarrow d$ and $a \rightarrow c \rightarrow d$. Thus, if we do not want to introduce new complications, we cannot rely on preclusion to decide conflicts. It seems that this would change the whole outlook on such diagrams. The interested reader will find more on the subject in [Ant97], [Ant99], [Ant05].

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 authors’ 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 and big subsets, familiar from preferential structures, see Definition 4.3 (page 29) .

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 (page 10) .

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 (page 11) , 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.

Diagram 2.2 LABEL: *Diagram Nixon-Diamond*

The Nixon Diamond



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 conclusions, see [MS91] for a detailed discussion. See also Section 2.3 (page 19) for more discussion on directly vs. indirectly sceptical approaches.

In more detail:

Preclusion

(+++*** Orig.: Preclusion)

LABEL: Section 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 $c \subseteq \dots \subseteq b$, where \subseteq stands here for soft inclusion), 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]).

A final remark: 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.

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

LABEL: Section IS-2.2

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, in particular in Section 2.3 (page 19)). 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 5.3 (page 36) 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 information in our sense. 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 5.4 (page 39) . (The fundamental difference between intersection of extensions and direct scepticism for defeasible inheritance was shown in [Sch93].) See also Section 2.3 (page 19) for more discussion.

We turn now to the announced variants as well as a finer distinction within the directly sceptical approach. Again, see also Section 2.3 (page 19) for more discussion.

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 bundle the decision of this problem with that for direct scepticism into a “plug-in” decision, which will be used in three approaches: the conceptual ideas, the inheritance algorithm, and the choice of the reference class for subset size (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) When we interpret arrows as causation (in the sense that $X \rightarrow Y$ expresses that condition X usually causes condition Y to result), 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 5.4 (page 39) 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. Thus, in Diagram 2.1 (page 10) , the comparison between the truth values “penguin” and “bird” is absolute, and does not depend on the point of view “Tweety”, as it can in total validity preclusion - if we continue to view preclusion as a comparison of information strength (or truth value). 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 5.4 (page 39) .

(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 IS-2.2)

LABEL: Definition IS-2.2

This is an informal definition of 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 radical 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)

(+++*** Orig.: The definition of m (i.e. of validity of paths))

LABEL: Section The definition of m (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)

LABEL: 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 $x \dots \rightarrow p$ is a generalized path, and $p \rightarrow q \in \Gamma$, then $x \dots \rightarrow p \rightarrow q$ is a generalized path,

if $x \dots \rightarrow p$ is a generalized path, and $p \not\rightarrow q \in \Gamma$, then $x \dots \rightarrow p \not\rightarrow q$ 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 this 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)

LABEL: Definition 2.4

LABEL: Definition ValidPaths

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

Let σ be a potential path.

- Case I:

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

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

- Case II:

σ is a compound potential path, $deg_{\Gamma,a}(\sigma) = n$, 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, informally, $\Gamma \models \sigma$ 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 (page 14) .

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)

LABEL: 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 popular 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 (page 10) , 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 validity preclusion, see also Section 2.3 (page 19) for more discussion, 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 $b \rightarrow d$ or $c \not\rightarrow d$ - 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)

LABEL: 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 $\sigma : x \cdots \not\rightarrow y$.

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

Fact 2.2

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

LABEL: Fact 2.2

The above definition and the informal one outlined in Definition 2.2 (page 14) correspond, when we consider valid positive paths as access to information and comparison of information strength as indicated at the beginning of Section 2.2 (page 13) and elaborated in Section 5.3 (page 36) .

Proof

(+++*** Orig.: Proof)

As Definition 2.2 (page 14) 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 (page 14) , and start from some x .

Case 1:

(+++*** Orig.: Case 1:)

LABEL: Section Case 1:

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

By comparison of strength via preclusion, 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:

(+++*** Orig.: Case 2:)

LABEL: Section 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 (page 14) .

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 (page 14) and II.1 (2) + (3) correspond. Note that variant (P2.1) of Definition 2.2 (page 14) would give a different result - which we could, of course, also imitate in a modified inheritance approach.

Case 3:

(+++*** Orig.: Case 3:)

LABEL: Section Case 3:

Other information.

Inheritance nets give no other information, as valid information is deduced only through valid paths by Definition 2.5 (page 17) . And we did not add any other information either in the approach in Definition 2.2 (page 14) . But as is obvious in Case 2, valid paths coincide in both cases.

Thus, both approaches are equivalent.

□

2.3 Review of other approaches and problems


LABEL: Section IS-2.3 LABEL: Section Review

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 2.3 (page 19) .

Diagram 2.3 LABEL: Diagram Up-Down-Chaining

The problem of downward chaining



Preclusions valid above (here at u) can be invalid at lower points (here at z), as part of the relevant information is no longer 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 also the remarks in Section 5.4 (page 39) .

Extension-based versus directly skeptical definitions

(+++*** Orig.: Extension-based versus directly skeptical definitions)

LABEL: Section 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 also the remark in Section 5.4 (page 39) .

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

On-path versus off-path preclusion

(+++*** Orig.: On-path versus off-path preclusion)

LABEL: Section 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 $a \rightarrow b \rightarrow d$.

Split validity versus total validity preclusion


(+++*** Orig.: Split validity versus total validity preclusion)

LABEL: Section 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 2.4 (page 20) , taken from [Sch97-2], the path $x \rightarrow w \rightarrow v$ is valid, so is $u \rightarrow x$, but not the whole preclusion path $u \rightarrow x \rightarrow w \rightarrow v$.

Diagram 2.4 LABEL: Diagram Split-Total-Preclusion

Split vs. total validity preclusion



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.

karl-search= End Is-Inh-Intro

karl-search= Start Is-Inh-Reac

3 Defeasible inheritance and reactive diagrams

LABEL: Section Def-Reac

Before we discuss the relationship in detail, we first summarize our algorithm.

3.1 Summary of our algorithm

We look for valid paths from x to y .

- (1) Direct arrows are valid paths.
- (2) Consider the set C of all direct predecessors of y , i.e. all c such that there is a direct link from c to y .
 - (2.1) Eliminate all those to which there is no valid positive path from x (found by induction), let the new set be $C' \subseteq C$.

If the existence of a valid path has not yet been decided, we have to wait.

- (2.2) Eliminate from C' all c such that there is $c' \in C'$ and a valid positive path from c' to c (found by induction) - unless the arrows from c and from c' to y are of the same type. Let the new set be $C'' \subseteq C'$ (this handles preclusion).

If the existence of such valid paths has not yet been decided, we have to wait.

(2.3) If the arrows from all elements of C'' to y have same polarity, we have a valid path from x to y , if not, we have an unresolved contradiction, and there is no such path.

Note that we were a bit sloppy here. It can be debated whether preclusion by some c' such that c and c' have the same type of arrow to y should be accepted. As we are basically only interested whether there is a valid path, but not in its details, this does not matter.

3.2 Overview

There are several ways to use reactive graphs to help us solve inheritance diagrams - but also to go beyond them.

- (1) We use them to record results found by adding suitable arrows to the graph.
- (2) We go deeper into the calculating mechanism, and do not only use new arrows as memory, but also for calculation.
- (3) We can put up “sign posts” to mark dead ends, in the following sense: If we memorize valid paths from x to y , then, anytime we are at a branching point u coming from x , trying to go to y , and there is valid path through an arrow leaving u , we can put up a sign post saying “no valid path from x to y through this arrow”.

Note that we have to state destination y (of course), but also outset, x : There might be a valid path from u to y , which may be precluded or contradicted by some path coming from x .

- (4) We can remember preclusion, in the following sense: If we found a valid positive path from a to b , and there are contradicting arrows from a and b to c , then we can create an arrow from a to the arrow from b to c . So, if, from x , we can reach both a and b , the arrow from b to c will be incapacitated.

Before we discuss the first three possibilities in detail we shortly discuss the more general picture (in rough outline).

- (1) Replacing labels by arrows and vice versa.

As we can switch arrows on and off, an arrow carries a binary value - even without any labels. So the idea is obvious:

If an arrow has 1 label with n possible values, we can replace it with n parallel arrows (i.e. same source and destination), where we switch exactly one of them on - this is the label.

Conversely, we can replace n parallel arrows without labels, where exactly one is active, by one arrow with n labels.

We can also code labels of a node x by an arrow $\alpha : x \rightarrow x$, which has the same labels.

- (2) Coding logical formulas and truth in a model

We take two arrows for each propositional variable, one stands for true, the other for false. Negation blocks the positive arrow, enables the negative one. Conjunction is solved by concatenation, disjunction by parallel paths. If a variable occurs more than once, we make copies, which are “switched” by the “master arrow”.

We come back to the first three ways to treat inheritance by reactive graphs, and also mention a way to go beyond usual inheritance.

3.3 Compilation and memorization

When we take a look at the algorithm deciding which potential paths are valid, we see that, with one exception, we only need the results already obtained, i.e. whether there is a valid positive/negative path from a to b , and not the actual paths themselves. (This is, of course, due to our split validity definition of preclusion.) The

exception is that preclusion works “in the upper part” with direct links. But this is a local problem: we only have to look at the direct predecessors of a node.

Consequently, we can do most of the calculation just once, in an induction of increasing “path distance”, memorize valid positive (the negative ones cannot be used) paths with special arrows which we activated once their validity is established, and work now as follows with the new arrows:

Suppose we want to know if there is a valid path from x to y . We look backwards at all predecessors b of y (using a simple backward pointer), and look whether there is a valid positive path from x to b , using the new arrows. We then look at all arrows going from such b to y . If they all agree (i.e. all are positive, or all are negative), we need not look further, and have the result. If they disagree, we have to look at possible comparisons by specificity. For this, we see whether there are new arrows between the b 's. All such b to which there is a new arrow from another b are out of consideration. If the remaining agree, we have a valid path (and activate a new arrow from x to y if the path is positive), if not, there is no such path. (Depending on the technical details of the induction, it can be useful to note this also by activating a special arrow.)

3.4 Executing the algorithm

Consider any two points x, y .

There can be no path, a positive potential path, a negative potential path, both potential paths, a valid positive path, a valid negative path (from x to y).

Once a valid path is found, we can forget potential paths, so we can code the possibilities by $\{*, p+, p-, p+ -, v+, v-\}$, in above order.

We saw that we can work either with labels, or with a multitude of parallel arrows, we choose the first possibility.

We create for each pair of nodes a new arrow, (x, y) , which we initialize with label $*$.

First, we look for potential paths.

If there is a direct link from x to y , we change the value $*$ of (x, y) directly to $v+$ or $v-$.

If (x, y') has value $p+$ or $p+ -$ or $v+$, and there is a direct link from y' to y , we change the value of (x, y) from $*$ to $p+$ or $p-$, depending on the link from y' to y , from $p+$ or $p-$ to $p+ -$ if adequate (we found both possibilities), and leave the value unchanged otherwise.

This determines all potential paths.

We make for each pair x, y at y a list of its predecessors, i.e. of all c s.t. there is a direct link from c to y . We do this for all x , so we can work in parallel. A list is, of course, a concatenation of suitable arrows.

Suppose we want to know if there is valid path from x to y .

First, there might be a direct link, and we are done.

Next, we look at the list of predecessors of y for x .

If one c in the list has value $*$, $p-$, $v-$ for (x, c) , it is eliminated from the list. If one c has $p+$ or $p+ -$, we have to wait. We do this until all (x, c) have $*$, $v+$ or $v-$, so those remaining in the list will have $v+$.

We look at all pairs in the list. While at least one (c, c') has $p+$ or $p+ -$, we have to wait. Finally, all pairs will have $*$, $p-$, $v+$ or $v-$. Eliminate all c' s.t. there is (c, c') with value $v+$ - unless the arrows from c and c' to y are of the same type.

Finally, we look at the list of the remaining predecessors, if they all have the same link to y , we set (x, y) to $v+$ or $v-$, otherwise to $*$.


All such operations can be done by suitable operations with arrows, but it is very lengthy and tedious to write down the details.

3.5 Signposts

Putting up signposts requires memorizing all valid paths, as leaving one valid path does not necessarily mean that there is no alternative valid path. The easiest thing to do this is probably to put up a warning post everywhere, and collect the wrong ones going backwards through the valid paths.

We illustrate this with Diagram 3.1 (page 24) :

Diagram 3.1 LABEL: *Diagram InherUniv*



There are the following potential paths from x to y : xcy , $xceb-y$, $xcebday$, xay .

The paths xc , xa , and $xceb$ are valid. The latter: xce is valid. $xceb$ is in competition with $xcg-b$ and $xceg-b$, but both are precluded by the arrow (valid path) eg .

y 's predecessors on those paths are a, b, c .

a and b are not comparable, as there is no valid path from b to a , as bda is contradicted by $bf-a$. None is more specific than the other one.

b and c are comparable, as the path ceb is valid, since $cg-b$ and $ceg-b$ are precluded by the valid path (arrow) eg . So c is more specific than b .

Thus, b is out of consideration, and we are left with a and c . They agree, so there is positive valid path from x to y , more precisely, one through a , one through c .

We have put STOP signs on the arrows ce and cg , as we cannot continue via them to y .

3.6 Beyond inheritance

But we can also go beyond usual inheritance networks.

Consider the following scenario:

- Museum airplanes usually will not fly, but usual airplanes will fly.
- Penguins *don't* fly, but birds do.
- Non-flying birds usually have fat wings.
- Non-flying aircraft usually are rusty.
- But it is not true that usually non-flying things are rusty and have fat wings.

We can model this with higher order arrows as follows:

- Penguins \rightarrow birds, museum airplanes \rightarrow airplanes, birds \rightarrow fly, airplanes \rightarrow fly.
- Penguins \nrightarrow fly, museum airplanes \nrightarrow fly.
- Flying objects \nrightarrow rusty, flying objects \nrightarrow have fat wings.
- We allow concatenation of two negative arrows:

Coming e.g. from penguins, we want to concatenate penguins \nrightarrow fly and fly \nrightarrow fat wings. Coming from museum aircraft, we want to concatenate museum aircraft \nrightarrow fly and fly \nrightarrow rusty.

We can enable this as follows: we introduce a new arrow $\alpha : (penguin \nrightarrow fly) \rightarrow (fly \nrightarrow fat\ wings)$, which, when traversing penguin \nrightarrow fly enables the algorithm to concatenate with the arrow it points to, using the rule " $- * - = +$ ", giving the result that penguins usually have fat wings.

See [Gab08c] for deeper discussion.

karl-search= End Is-Inh-Reac

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

karl-search= Start Is-Small-Intro

LABEL: Section Is-Small-Intro

We now continue to prepare the ground for our conceptual analysis of inheritance nets. We need some technical results and we devote this section just to this end.

Definition 4.1

(+++*** Orig. No.: Definition Log-Base)

LABEL: Definition Log-Base

We work here in a classical propositional language \mathcal{L} , a theory T will be an arbitrary set of formulas. Formulas will often be named ϕ, ψ , etc., theories T, S , etc.

$v(\mathcal{L})$ will be the set of propositional variables of \mathcal{L} .

$M_{\mathcal{L}}$ will be the set of (classical) models for \mathcal{L} , $M(T)$ or M_T is the set of models of T , likewise $M(\phi)$ for a formula ϕ .

$\mathbf{D}_{\mathcal{L}} := \{M(T) : T \text{ a theory in } \mathcal{L}\}$, the set of definable model sets.

Note that, in classical propositional logic, $\emptyset, M_{\mathcal{L}} \in \mathbf{D}_{\mathcal{L}}$, $\mathbf{D}_{\mathcal{L}}$ contains singletons, is closed under arbitrary intersections and finite unions.

An operation $f : \mathcal{Y} \rightarrow \mathcal{P}(M_{\mathcal{L}})$ for $\mathcal{Y} \subseteq \mathcal{P}(M_{\mathcal{L}})$ is called definability preserving, (*dp*) or (*mdp*) in short, iff for all $X \in \mathbf{D}_{\mathcal{L}} \cap \mathcal{Y}$ $f(X) \in \mathbf{D}_{\mathcal{L}}$.

We will also use (*mdp*) for binary functions $f : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{P}(M_{\mathcal{L}})$ - as needed for theory revision - with the obvious meaning.

\vdash will be classical derivability, and

$\overline{T} := \{\phi : T \vdash \phi\}$, the closure of T under \vdash .

$Con(\cdot)$ will stand for classical consistency, so $Con(\phi)$ will mean that ϕ is classical consistent, likewise for $Con(T)$. $Con(T, T')$ will stand for $Con(T \cup T')$, etc.

Given a consequence relation \vdash , we define

$\overline{\overline{T}} := \{\phi : T \vdash \phi\}$.

(There is no fear of confusion with $\overline{\overline{T}}$, as it just is not useful to close twice under classical logic.)

$T \vee T' := \{\phi \vee \phi' : \phi \in T, \phi' \in T'\}$.

If $X \subseteq M_{\mathcal{L}}$, then $Th(X) := \{\phi : X \models \phi\}$, likewise for $Th(m)$, $m \in M_{\mathcal{L}}$. (\models will usually be classical validity.)

Definition 4.2

(+++*** Orig. No.: Definition Log-Cond)

LABEL: Definition Log-Cond

We introduce here formally a list of properties of set functions on the algebraic side, and their corresponding logical rules on the other side.

Recall that $\overline{T} := \{\phi : T \vdash \phi\}$, $\overline{\overline{T}} := \{\phi : T \vdash \phi\}$, where \vdash is classical consequence, and \vdash any other consequence.

We show, wherever adequate, in parallel the formula version in the left column, the theory version in the middle column, and the semantical or algebraic counterpart in the right column. The algebraic counterpart gives conditions for a function $f : \mathcal{Y} \rightarrow \mathcal{P}(U)$, where U is some set, and $\mathcal{Y} \subseteq \mathcal{P}(U)$.

Precise connections between the columns are given in Proposition 4.6 (page 33)

When the formula version is not commonly used, we omit it, as we normally work only with the theory version.

A and B in the right hand side column stand for $M(\phi)$ for some formula ϕ , whereas X, Y stand for $M(T)$ for some theory T .

Basics		
(AND) $\phi \vdash \psi, \phi \vdash \psi' \Rightarrow$ $\phi \vdash \psi \wedge \psi'$	(AND) $T \vdash \psi, T \vdash \psi' \Rightarrow$ $T \vdash \psi \wedge \psi'$	Closure under finite intersection
(OR) $\phi \vdash \psi, \phi' \vdash \psi \Rightarrow$ $\phi \vee \phi' \vdash \psi$	(OR) $\overline{\overline{T}} \cap \overline{\overline{T'}} \subseteq \overline{\overline{T \vee T'}}$	(μ OR) $f(X \cup Y) \subseteq f(X) \cup f(Y)$
(wOR) $\phi \vdash \psi, \phi' \vdash \psi \Rightarrow$ $\phi \vee \phi' \vdash \psi$	(wOR) $\overline{\overline{T}} \cap \overline{\overline{T'}} \subseteq \overline{\overline{T \vee T'}}$	(μ wOR) $f(X \cup Y) \subseteq f(X) \cup Y$
(disjOR) $\phi \vdash \neg \phi', \phi \vdash \psi,$ $\phi' \vdash \psi \Rightarrow \phi \vee \phi' \vdash \psi$	(disjOR) $\neg \text{Con}(T \cup T') \Rightarrow$ $\overline{\overline{T}} \cap \overline{\overline{T'}} \subseteq \overline{\overline{T \vee T'}}$	(μ disjOR) $X \cap Y = \emptyset \Rightarrow$ $f(X \cup Y) \subseteq f(X) \cup f(Y)$
(LLE) Left Logical Equivalence $\vdash \phi \leftrightarrow \phi', \phi \vdash \psi \Rightarrow$ $\phi' \vdash \psi$	(LLE) $\overline{\overline{T}} = \overline{\overline{T'}} \Rightarrow \overline{\overline{T}} = \overline{\overline{T'}}$	trivially true
(RW) Right Weakening $\phi \vdash \psi, \vdash \psi \rightarrow \psi' \Rightarrow$ $\phi \vdash \psi'$	(RW) $T \vdash \psi, \vdash \psi \rightarrow \psi' \Rightarrow$ $T \vdash \psi'$	upward closure
(CCL) Classical Closure	(CCL) $\overline{\overline{T}}$ is classically closed	trivially true
(SC) Supraclassicality $\phi \vdash \psi \Rightarrow \phi \vdash \psi$	(SC) $\overline{\overline{T}} \subseteq \overline{\overline{T}}$	($\mu \subseteq$) $f(X) \subseteq X$
(REF) Reflexivity $T \cup \{\alpha\} \vdash \alpha$		
(CP) Consistency Preservation $\phi \vdash \perp \Rightarrow \phi \vdash \perp$	(CP) $T \vdash \perp \Rightarrow T \vdash \perp$	($\mu \emptyset$) $f(X) = \emptyset \Rightarrow X = \emptyset$
		($\mu \emptyset \text{fin}$) $X \neq \emptyset \Rightarrow f(X) \neq \emptyset$ for finite X
$\overline{\overline{\phi \wedge \phi'}} \subseteq \overline{\overline{\phi \cup \{\phi'\}}}$	(PR) $\overline{\overline{T \cup T'}} \subseteq \overline{\overline{T \cup T'}}$	(μ PR) $X \subseteq Y \Rightarrow$ $f(Y) \cap X \subseteq f(X)$ (μ PR') $f(X) \cap Y \subseteq f(X \cap Y)$
(CUT) $T \vdash \alpha; T \cup \{\alpha\} \vdash \beta \Rightarrow$ $T \vdash \beta$	(CUT) $T \subseteq \overline{\overline{T'}} \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T'}} \subseteq \overline{\overline{T}}$	(μ CUT) $f(X) \subseteq Y \subseteq X \Rightarrow$ $f(X) \subseteq f(Y)$

Cumulativity		
(<i>CM</i>) Cautious Monotony $\phi \sim \psi, \phi \sim \psi' \Rightarrow$ $\phi \wedge \psi \sim \psi'$	(<i>CM</i>) $T \subseteq \overline{T'} \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T}} \subseteq \overline{\overline{T'}}$	(μ <i>CM</i>) $f(X) \subseteq Y \subseteq X \Rightarrow$ $f(Y) \subseteq f(X)$
or (<i>ResM</i>) Restricted Monotony $T \sim \alpha, \beta \Rightarrow T \cup \{\alpha\} \sim \beta$		(μ <i>ResM</i>) $f(X) \subseteq A \cap B \Rightarrow f(X \cap A) \subseteq B$
(<i>CUM</i>) Cumulativity $\phi \sim \psi \Rightarrow$ $(\phi \sim \psi' \Leftrightarrow \phi \wedge \psi \sim \psi')$	(<i>CUM</i>) $T \subseteq \overline{T'} \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T}} = \overline{\overline{T'}}$	(μ <i>CUM</i>) $f(X) \subseteq Y \subseteq X \Rightarrow$ $f(Y) = f(X)$
	($\subseteq \supseteq$) $T \subseteq \overline{\overline{T'}}, T' \subseteq \overline{\overline{T}} \Rightarrow$ $\overline{\overline{T'}} = \overline{\overline{T}}$	($\mu \subseteq \supseteq$) $f(X) \subseteq Y, f(Y) \subseteq X \Rightarrow$ $f(X) = f(Y)$
Rationality		
(<i>RatM</i>) Rational Monotony $\phi \sim \psi, \phi \not\sim \neg\psi' \Rightarrow$ $\phi \wedge \psi' \sim \psi$	(<i>RatM</i>) $Con(T \cup \overline{\overline{T'}}), T \vdash T' \Rightarrow$ $\overline{\overline{T}} \supseteq \overline{\overline{T'}} \cup T$	(μ <i>RatM</i>) $X \subseteq Y, X \cap f(Y) \neq \emptyset \Rightarrow$ $f(X) \subseteq f(Y) \cap X$
	(<i>RatM</i> =) $Con(T \cup \overline{\overline{T'}}), T \vdash T' \Rightarrow$ $\overline{\overline{T}} = \overline{\overline{T'}} \cup T$	(μ =) $X \subseteq Y, X \cap f(Y) \neq \emptyset \Rightarrow$ $f(X) = f(Y) \cap X$
	(<i>Log</i> =') $Con(\overline{\overline{T'}} \cup T) \Rightarrow$ $\overline{\overline{T \cup T'}} = \overline{\overline{T'}} \cup T$	(μ =') $f(Y) \cap X \neq \emptyset \Rightarrow$ $f(Y \cap X) = f(Y) \cap X$
	(<i>Log</i>) $\overline{\overline{T \vee T'}}$ is one of $\overline{\overline{T}},$ or $\overline{\overline{T'}},$ or $\overline{\overline{T}} \cap \overline{\overline{T'}}$ (by (CCL))	(μ) $f(X \cup Y)$ is one of $f(X), f(Y)$ or $f(X) \cup f(Y)$
	(<i>Log</i> \cup) $Con(\overline{\overline{T'}} \cup T), \neg Con(\overline{\overline{T'}} \cup \overline{\overline{T}}) \Rightarrow$ $\neg Con(\overline{\overline{T \vee T'}} \cup T')$	($\mu \cup$) $f(Y) \cap (X - f(X)) \neq \emptyset \Rightarrow$ $f(X \cup Y) \cap Y = \emptyset$
	(<i>Log</i> \cup') $Con(\overline{\overline{T'}} \cup T), \neg Con(\overline{\overline{T'}} \cup \overline{\overline{T}}) \Rightarrow$ $\overline{\overline{T \vee T'}} = \overline{\overline{T}}$	($\mu \cup'$) $f(Y) \cap (X - f(X)) \neq \emptyset \Rightarrow$ $f(X \cup Y) = f(X)$
		($\mu \in$) $a \in X - f(X) \Rightarrow$ $\exists b \in X. a \notin f(\{a, b\})$

(*PR*) is also called infinite conditionalization - we choose the name for its central role for preferential structures (*PR*) or (μ *PR*).

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

Roughly: Smooth preferential structures generate logics satisfying system *P*, ranked structures logics satisfying system *R*.

A logic satisfying (*REF*), (*ResM*), and (*CUT*) is called a consequence relation.

(*LLE*) and (*CCL*) will hold automatically, whenever we work with model sets.

(*AND*) is obviously closely related to filters, and corresponds to closure under finite intersections. (*RW*) corresponds to upward closure of filters.

More precisely, validity of both depend on the definition, and the direction we consider.

Given f and ($\mu \subseteq$), $f(X) \subseteq X$ generates a principal filter: $\{X' \subseteq X : f(X) \subseteq X'\}$, with the definition: If $X = M(T)$, then $T \sim \phi$ iff $f(X) \subseteq M(\phi)$. Validity of (*AND*) and (*RW*) are then trivial.

Conversely, we can define for $X = M(T)$

$$\mathcal{X} := \{X' \subseteq X : \exists \phi (X' = X \cap M(\phi) \text{ and } T \sim \phi)\}.$$

(*AND*) then makes \mathcal{X} closed under finite intersections, (*RW*) makes \mathcal{X} upward closed. This is in the infinite case usually not yet a filter, as not all subsets of X need to be definable this way. In this case, we complete \mathcal{X} by adding all X'' such that there is $X' \subseteq X'' \subseteq X$, $X' \in \mathcal{X}$.

Alternatively, we can define

$$\mathcal{X} := \{X' \subseteq X : \bigcap \{X \cap M(\phi) : T \sim \phi\} \subseteq X'\}.$$

(*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; the problem is addressed by the limit versions.

(*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 more 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.

Definition 4.3

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

LABEL: Definition Filter

A filter is an abstract notion of size, elements of a filter $\mathcal{F}(X)$ on X are called big subsets of X , their complements are called small, and the rest have medium size. The dual applies to ideals $\mathcal{I}(X)$, this is justified by the trivial fact that $\{X - A : A \in \mathcal{F}(X)\}$ is an ideal iff $\mathcal{F}(X)$ is a filter.

In both definitions, the first two conditions (i.e. (*FAll*), (*I*∅), and (*F*↑), (*I*↓)) should hold if the notions shall have anything to do with usual intuition, and there are reasons to consider only the weaker, less idealistic, version of the third.

At the same time, we introduce - in rough parallel - coherence conditions which describe what might happen when we change the reference or base set X . (*R*↑) is very natural, (*R*↓) is more daring, and (*R*↓↓) even more so. (*R*∪*disj*) is a cautious combination of (*R*↑) and (*R*∪), as we avoid using the same big set several times in comparison, so (*R*∪) is used more cautiously here. See Remark 4.1 (page 30) for more details.

Finally, we give a generalized first order quantifier corresponding to a (weak) filter. The precise connection is formulated in Definition 4.4 (page 31) , Definition 4.5 (page 32) , Definition 4.6 (page 32) , and Proposition 4.4 (page 32) , respectively their relativized versions.

Fix now a base set $X \neq \emptyset$.

A (weak) filter on or over X is a set $\mathcal{F}(X) \subseteq \mathcal{P}(X)$, s.t. (*FAll*), (*F*↑), (*F*∩) ((*FAll*), (*F*↑), (*F*∩')) respectively hold.

A filter is called a principal filter iff there is $X' \subseteq X$ s.t. $\mathcal{F} = \{A : X' \subseteq A \subseteq X\}$.

A filter is called an ultrafilter iff for all $X' \subseteq X$ $X' \in \mathcal{F}(X)$ or $X - X' \in \mathcal{F}(X)$.

A (weak) ideal on or over X is a set $\mathcal{I}(X) \subseteq \mathcal{P}(X)$, s.t. (*I*∅), (*I*↓), (*I*∪) ((*I*∅), (*I*↓), (*I*∪')) respectively hold.

Finally, we set $\mathcal{M}(X) := \{A \subseteq X : A \notin \mathcal{I}(X), A \notin \mathcal{F}(X)\}$, the “medium size” sets, and $\mathcal{M}^+(X) := \mathcal{M}(X) \cup \mathcal{F}(X)$, $\mathcal{M}^+(X)$ is the set of subsets of X , which are not small, i.e. have medium or large size.

For (*R*↓) and (*R*↓↓) (–) is assumed in the following table.

Optimum			
$(FAll)$ $X \in \mathcal{F}(X)$	$(I\emptyset)$ $\emptyset \in \mathcal{I}(X)$		$\forall x\phi(x) \rightarrow \nabla x\phi(x)$
Improvement			
$(F\uparrow)$ $A \subseteq B \subseteq X,$ $A \in \mathcal{F}(X) \Rightarrow$ $B \in \mathcal{F}(X)$	$(I\downarrow)$ $A \subseteq B \subseteq X,$ $B \in \mathcal{I}(X) \Rightarrow$ $A \in \mathcal{I}(X)$	$(R\uparrow)$ $X \subseteq Y \Rightarrow \mathcal{I}(X) \subseteq \mathcal{I}(Y)$	$\nabla x\phi(x) \wedge$ $\forall x(\phi(x) \rightarrow \psi(x)) \rightarrow$ $\nabla x\psi(x)$
Adding small sets			
$(F\cap)$ $A, B \in \mathcal{F}(X) \Rightarrow$ $A \cap B \in \mathcal{F}(X)$	$(I\cup)$ $A, B \in \mathcal{I}(X) \Rightarrow$ $A \cup B \in \mathcal{I}(X)$	$(R\downarrow)$ $A, B \in \mathcal{I}(X) \Rightarrow$ $A - B \in \mathcal{I}(X-B)$ or: $A \in \mathcal{F}(X), B \in \mathcal{I}(X) \Rightarrow$ $A - B \in \mathcal{F}(X-B)$	$\nabla x\phi(x) \wedge \nabla x\psi(x) \rightarrow$ $\nabla x(\phi(x) \wedge \psi(x))$
Cautious addition			
$(F\cap')$ $A, B \in \mathcal{F}(X) \Rightarrow$ $A \cap B \neq \emptyset.$	$(I\cup')$ $A, B \in \mathcal{I}(X) \Rightarrow$ $A \cup B \neq X.$	$(R\cup disj)$ $A \in \mathcal{I}(X), B \in \mathcal{I}(Y), X \cap Y = \emptyset \Rightarrow$ $A \cup B \in \mathcal{I}(X \cup Y)$	$\nabla x\phi(x) \rightarrow \nabla x\neg\phi(x)$ and $\nabla x\phi(x) \rightarrow \exists x\phi(x)$
Bold addition			
Ultrafilter	(Dual of) Ultrafilter	$(R\downarrow\downarrow)$ $A \in \mathcal{I}(X), B \notin \mathcal{F}(X) \Rightarrow$ $A - B \in \mathcal{I}(X - B)$ or: $A \in \mathcal{F}(X), B \notin \mathcal{F}(X) \Rightarrow$ $A - B \in \mathcal{F}(X - B)$ or: $A \in \mathcal{M}^+(X), X \in \mathcal{M}^+(Y) \Rightarrow$ $A \in \mathcal{M}^+(Y)$ - Transitivity of \mathcal{M}^+	$\neg\nabla x\phi(x) \rightarrow \nabla x\neg\phi(x)$

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.

Seen more abstractly, set properties like e.g. $(R\uparrow)$ allow the transfer of big (or small) 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.

Remark 4.1

(+++*** Orig. No.: Remark Ref-Class)

LABEL: Remark Ref-Class

$(R\uparrow)$ corresponds to $(I\downarrow)$ and $(F\uparrow)$: If A is small in $X \subseteq Y$, then it will a fortiori be small in the bigger Y .

$(R\downarrow)$ says that diminishing base sets by a small amount will keep small subsets small. This goes in the wrong direction, so we have to be careful. We cannot diminish arbitrarily, e.g., if A is a small subset of B , A should not be a small subset of $B - (B - A) = A$. It still seems quite safe, if “small” is a robust notion, i.e. defined in an abstract way, and not anew for each set, and, if “small” is sufficiently far from “big”, as, for example in a filter.

There is, however, an important conceptual distinction to make here. Filters express “size” in an abstract way, in the context of nonmonotonic logics, $\alpha \sim \beta$ iff the set of $\alpha \wedge \neg\beta$ is small in α . But here, we were interested in “small” changes in the reference set X (or α in our example). So we have two quite different uses of “size”,

one for nonmonotonic logics, abstractly expressed by a filter, the other for coherence conditions. It is possible, but not necessary, to consider both essentially the same notions. But we should not forget that we have two conceptually different uses of size here.

$(R \downarrow \downarrow)$ is obviously a stronger variant of $(R \downarrow)$.

It and its strength is perhaps best understood as transitivity of the relation $xSy : \Leftrightarrow x \in \mathcal{M}^+(y)$.

Now, (in comparison to $(R \downarrow)$) A' can be a medium size subset of B . As a matter of fact, $(R \downarrow \downarrow)$ is a very big strengthening of $(R \downarrow)$: Consider a principal filter $\mathcal{F} := \{X \subseteq B : B' \subseteq X\}$, $b \in B'$. Then $\{b\}$ has at least medium size, so any small set $A \subseteq B$ is smaller than $\{b\}$ - and this is, of course, just rankedness. If we only have $(R \downarrow)$, then we need the whole generating set B' to see that A is small. This is the strong substitution property of rankedness: any b as above will show that A is small.

The more we see size as an abstract notion, and the more we see “small” different from “big” (or “medium”), the more we can go from one base set to another and find the same sizes - the more we have coherence when we reason with small and big subsets. $(R \downarrow)$ works with iterated use of “small”, just as do filters, but not weak filters. So it is not surprising that weak filters and $(R \downarrow)$ do not cooperate well: Let A, B, C be small subsets of X - pairwise disjoint, and $A \cup B \cup C = X$, this is possible. By $(R \downarrow)$ B and C will be small in $X - A$, so again by $(R \downarrow)$ C will be small in $(X - A) - B = C$, but this is absurd.

If we think that filters are too strong, but we still want some coherence, i.e. abstract size, we can consider $(R \cup disj)$: If A is a small subset of B , and A' of B' , and B and B' are disjoint, then $A \cup A'$ is a small subset of $B \cup B'$. It expresses a uniform approach to size, or distributivity, if you like. It holds, e.g. when we consider a set to be small iff it is smaller than a certain fraction. The important point is here that by disjointness, the big subsets do not get “used up”. (This property generalizes in a straightforward way to the infinite case.)

Fact 4.2

(+++*** Orig. No.: Fact R-down)

LABEL: Fact R-down

The two versions of $(R \downarrow)$ and the three versions of $(R \downarrow \downarrow)$ are each equivalent. For the third version of $(R \downarrow \downarrow)$ we use $(I \downarrow)$.

Proposition 4.3

(+++*** Orig. No.: Proposition Ref-Class-Mu)

LABEL: Proposition Ref-Class-Mu

If $f(X)$ is the smallest A s.t. $A \in \mathcal{F}(X)$, then, given the property on the left, the one on the right follows.

Conversely, when we define $\mathcal{F}(X) := \{X' : f(X) \subseteq X' \subseteq X\}$, given the property on the right, the one on the left follows. For this direction, we assume that we can use the full powerset of some base set U - as is the case for the model sets of a finite language. This is perhaps not too bold, as we mainly want to stress here the intuitive connections, without putting too much weight on definability questions.

(1.1)	$(R \uparrow)$	\Rightarrow	(μwOR)
(1.2)		\Leftarrow	
(2.1)	$(R \uparrow) + (I \cup)$	\Rightarrow	(μOR)
(2.2)		\Leftarrow	
(3.1)	$(R \uparrow) + (I \cup)$	\Rightarrow	(μPR)
(3.2)		\Leftarrow	
(4.1)	$(R \cup disj)$	\Rightarrow	$(\mu disjOR)$
(4.2)		\Leftarrow	
(5.1)	$(R \downarrow)$	\Rightarrow	(μCM)
(5.2)		\Leftarrow	
(6.1)	$(R \downarrow \downarrow)$	\Rightarrow	$(\mu RatM)$
(6.2)		\Leftarrow	

Definition 4.4

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

LABEL: Definition Nabla

Augment the language of first order logic by the new quantifier: If ϕ and ψ are formulas, then so are $\nabla x\phi(x)$, $\nabla x\phi(x) : \psi(x)$, for any variable x . The ∇ -versions are the restricted variants. We call any formula of \mathcal{L} , possibly containing ∇ a ∇ - \mathcal{L} -formula.

Definition 4.5

(+++*** Orig. No.: Definition N-Model)

LABEL: Definition N-Model

(\mathcal{N} -Model)

Let \mathcal{L} be a first order language, and M be a \mathcal{L} -structure. Let $\mathcal{N}(M)$ be a weak filter, or \mathcal{N} -system - \mathcal{N} for normal - over M . Define $\langle M, \mathcal{N}(M) \rangle \models \phi$ for any ∇ - \mathcal{L} -formula inductively as usual, with one additional induction step:

$\langle M, \mathcal{N}(M) \rangle \models \nabla x\phi(x)$ iff there is $A \in \mathcal{N}(M)$ s.t. $\forall a \in A (\langle M, \mathcal{N}(M) \rangle \models \phi[a])$.

Definition 4.6

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

LABEL: Definition NablaAxioms

Let any axiomatization of predicate calculus be given. Augment this with the axiom schemata

- (1) $\nabla x\phi(x) \wedge \forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla x\psi(x)$,
 - (2) $\nabla x\phi(x) \rightarrow \neg\nabla x\neg\phi(x)$,
 - (3) $\forall x\phi(x) \rightarrow \nabla x\phi(x)$ and $\nabla x\phi(x) \rightarrow \exists x\phi(x)$,
 - (4) $\nabla x\phi(x) \leftrightarrow \nabla y\phi(y)$ if x does not occur free in $\phi(y)$ and y does not occur free in $\phi(x)$.
- (for all ϕ, ψ).

Proposition 4.4

(+++*** Orig. No.: Proposition NablaRepr)

LABEL: Proposition NablaRepr

The axioms given in Definition 4.6 (page 32) are sound and complete for the semantics of Definition 4.5 (page 32)

See [Sch95-1] or [Sch04].

Definition 4.7

(+++*** Orig. No.: Definition Nabla-System)

LABEL: Definition Nabla-System

Call $\mathcal{N}^+(M) = \langle \mathcal{N}(N) : N \subseteq M \rangle$ a \mathcal{N}^+ -system or system of weak filters over M iff for each $N \subseteq M$ $\mathcal{N}(N)$ is a weak filter or \mathcal{N} -system over N . (It suffices to consider the definable subsets of M .)

Definition 4.8

(+++*** Orig. No.: Definition N-Model-System)

LABEL: Definition N-Model-System

Let \mathcal{L} be a first order language, and M a \mathcal{L} -structure. Let $\mathcal{N}^+(M)$ be a \mathcal{N}^+ -system over M .

Define $\langle M, \mathcal{N}^+(M) \rangle \models \phi$ for any formula inductively as usual, with the additional induction steps:

1. $\langle M, \mathcal{N}^+(M) \rangle \models \nabla x\phi(x)$ iff there is $A \in \mathcal{N}(M)$ s.t. $\forall a \in A (\langle M, \mathcal{N}^+(M) \rangle \models \phi[a])$,
2. $\langle M, \mathcal{N}^+(M) \rangle \models \nabla x\phi(x) : \psi(x)$ iff there is $A \in \mathcal{N}(\{x : \langle M, \mathcal{N}^+(M) \rangle \models \phi(x)\})$ s.t. $\forall a \in A (\langle M, \mathcal{N}^+(M) \rangle \models \psi[a])$.

Definition 4.9

(+++*** Orig. No.: Definition NablaAxioms-System)

LABEL: Definition NablaAxioms-System

Extend the logic of first order predicate calculus by adding the axiom schemata

- (1) a. $\nabla x\phi(x) \leftrightarrow \nabla x(x = x) : \phi(x)$, b. $\forall x(\sigma(x) \leftrightarrow \tau(x)) \wedge \nabla x\sigma(x) : \phi(x) \rightarrow \nabla x\tau(x) : \phi(x)$,
 - (2) $\nabla x\phi(x) : \psi(x) \wedge \forall x(\phi(x) \wedge \psi(x) \rightarrow \vartheta(x)) \rightarrow \nabla x\phi(x) : \vartheta(x)$,
 - (3) $\exists x\phi(x) \wedge \nabla x\phi(x) : \psi(x) \rightarrow \neg\nabla x\phi(x) : \neg\psi(x)$,
 - (4) $\forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla x\phi(x) : \psi(x)$ and $\nabla x\phi(x) : \psi(x) \rightarrow [\exists x\phi(x) \rightarrow \exists x(\phi(x) \wedge \psi(x))]$,
 - (5) $\nabla x\phi(x) : \psi(x) \leftrightarrow \nabla y\phi(y) : \psi(y)$ (under the usual caveat for substitution).
- (for all $\phi, \psi, \vartheta, \sigma, \tau$).

Proposition 4.5

(+++*** Orig. No.: Proposition NablaRepr-System)

LABEL: Proposition NablaRepr-System

The axioms of Definition 4.9 (page 33) are sound and complete for the \mathcal{N}^+ – semantics of ∇ as defined in Definition 4.8 (page 32) .

See [Sch95-1] or [Sch04].

karl-search= End Is-Small-Intro

Proposition 4.6

(+++*** Orig. No.: Proposition Alg-Log)

LABEL: Proposition Alg-Log

The following table is to be read as follows:

Let a logic \sim satisfy (LLE) and (CCL), and define a function $f : \mathbf{D}_{\mathcal{L}} \rightarrow \mathbf{D}_{\mathcal{L}}$ by $f(M(T)) := M(\overline{\overline{T}})$. Then f is well defined, satisfies (μdp), and $\overline{\overline{T}} = Th(f(M(T)))$.

If \sim satisfies a rule in the left hand side, then - provided the additional properties noted in the middle for \Rightarrow hold, too - f will satisfy the property in the right hand side.

Conversely, if $f : \mathcal{Y} \rightarrow \mathcal{P}(M_{\mathcal{L}})$ is a function, with $\mathbf{D}_{\mathcal{L}} \subseteq \mathcal{Y}$, and we define a logic \sim by $\overline{\overline{T}} := Th(f(M(T)))$, then \sim satisfies (LLE) and (CCL). If f satisfies (μdp), then $f(M(T)) = M(\overline{\overline{T}})$.

If f satisfies a property in the right hand side, then - provided the additional properties noted in the middle for \Leftarrow hold, too - \sim will satisfy the property in the left hand side.

If “formula” is noted in the table, this means that, if one of the theories (the one named the same way in Definition 4.2 (page 26)) is equivalent to a formula, we do not need (μdp).

Basics			
(1.1)	(OR)	\Rightarrow	(μOR)
(1.2)		\Leftarrow	
(2.1)	$(disjOR)$	\Rightarrow	$(\mu disjOR)$
(2.2)		\Leftarrow	
(3.1)	(wOR)	\Rightarrow	(μwOR)
(3.2)		\Leftarrow	
(4.1)	(SC)	\Rightarrow	$(\mu \subseteq)$
(4.2)		\Leftarrow	
(5.1)	(CP)	\Rightarrow	$(\mu \emptyset)$
(5.2)		\Leftarrow	
(6.1)	(PR)	\Rightarrow	(μPR)
(6.2)		$\Leftarrow (\mu dp) + (\mu \subseteq)$	
(6.3)		$\not\Leftarrow$ without (μdp)	
(6.4)		$\Leftarrow (\mu \subseteq)$ T' a formula	
(6.5)	(PR)	\Leftarrow T' a formula	$(\mu PR')$
(7.1)	(CUT)	\Rightarrow	(μCUT)
(7.2)		\Leftarrow	
Cumulativity			
(8.1)	(CM)	\Rightarrow	(μCM)
(8.2)		\Leftarrow	
(9.1)	$(ResM)$	\Rightarrow	$(\mu ResM)$
(9.2)		\Leftarrow	
(10.1)	$(\subseteq \supseteq)$	\Rightarrow	$(\mu \subseteq \supseteq)$
(10.2)		\Leftarrow	
(11.1)	(CUM)	\Rightarrow	(μCUM)
(11.2)		\Leftarrow	
Rationality			
(12.1)	$(RatM)$	\Rightarrow	$(\mu RatM)$
(12.2)		$\Leftarrow (\mu dp)$	
(12.3)		$\not\Leftarrow$ without (μdp)	
(12.4)		\Leftarrow T a formula	
(13.1)	$(RatM =)$	\Rightarrow	$(\mu =)$
(13.2)		$\Leftarrow (\mu dp)$	
(13.3)		$\not\Leftarrow$ without (μdp)	
(13.4)		\Leftarrow T a formula	
(14.1)	$(Log =')$	\Rightarrow	$(\mu =')$
(14.2)		$\Leftarrow (\mu dp)$	
(14.3)		$\not\Leftarrow$ without (μdp)	
(14.4)		$\Leftarrow T$ a formula	
(15.1)	$(Log \parallel)$	\Rightarrow	$(\mu \parallel)$
(15.2)		\Leftarrow	
(16.1)	$(Log \cup)$	$\Rightarrow (\mu \subseteq) + (\mu =)$	$(\mu \cup)$
(16.2)		$\Leftarrow (\mu dp)$	
(16.3)		$\not\Leftarrow$ without (μdp)	
(17.1)	$(Log \cup')$	$\Rightarrow (\mu \subseteq) + (\mu =)$	$(\mu \cup')$
(17.2)		$\Leftarrow (\mu dp)$	
(17.3)		$\not\Leftarrow$ without (μdp)	

5 Interpretations

karl-search= Start Is-Interpr

LABEL: Section Is-Interpr

5.1 Introduction

LABEL: Section 4.1

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 6 (page 41), where an interpretation in terms of small sets will be tried nonetheless, and its limitations explored.

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 of deontic and of 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.

5.2 Informal comparison of inheritance with the systems P and R

LABEL: Section 4.2

The main issues

(+++*** Orig.: The main issues)

LABEL: Section The main issues

In the authors' 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

(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 in Fact 6.1 (page 42) that b 's are non- c 's in system P - see Definition 4.2 (page 26)). 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.

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 4.2 (page 26)) differ on information transfer. In both cases, transfer is based on the same notion of smallness, which describes ideal situations. But, as said in Remark 4.1 (page 30), this is conceptually very different from the use of 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 6 (page 41) .

Note that Reiter Defaults behave much more like inheritance nets than like preferential logics.

5.3 Inheritance as information transfer

LABEL: Section 4.3

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. 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, arrows, and valid paths have multiple functions, 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 authors' 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:

(+++*** Orig.: (1) Information:)

LABEL: Section (1) Information:

Direct positive or negative arrows 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:

(+++*** Orig.: (2) Information sources and flow:)

LABEL: Section (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 can 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:

(+++*** Orig.: (3) Information strength:)

LABEL: Section (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 the much more complicated idea of relative comparison.

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?

IS-Discussion:

(+++*** Orig.: IS-Discussion:)

LABEL: Section IS-Discussion:

Considering direct arrows as information meets probably with little objection.

The conclusion of a valid path (e.g. if $\sigma : a \dots \rightarrow b$ 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 just 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 their sources, 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:

(+++*** Orig.: Illustration:)

LABEL: Section 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 (page 10), 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 we are interested in.

Note that negative links and (valid) paths have much less function in our picture than positive links and valid paths. In a way, this asymmetry is not surprising, as there are no negative nodes (which would correspond to something like the set complement or negation). To summarize: 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 in the sense of direct links. This distinction is very important, and corresponds to the different treatment of direct arrows and valid paths in inheritance, as it appears e.g. in the definition of preclusion. 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.

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

We should perhaps note in this context a connection to an area currently en vogue: the problem of trust, especially in the context of web information. We can see our truth values as the degrees of trust we put into information coming from this node. And, we not only use, but also reason about them.

Further comments:

(+++*** Orig.: Further comments:)

LABEL: Section 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 authors make 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 C complementation): If, for some X , we also have a node for CX , then we should have $X \not\rightarrow CX$ and $CX \not\rightarrow X$, thus a cycle, and arrows from Y to X should be accompanied by their opposite to CX , etc.

We translate the analysis and decision of Definition 2.2 (page 14) 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 (page 17) , (4), and turn this into a formal system.

Fix a diagram Γ , and do an induction as in Definition 2.2 (page 14) .

Definition 5.1

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

LABEL: 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.

(We introduce this notation to point informally to our idea of information strength, and beyond, to logical systems with varying strengths of implications.)

(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 that we cannot make $(a \Rightarrow b \text{ and } b \Rightarrow c \text{ imply } a \Rightarrow_b c)$ a rule, as this would make concatenation of two composed paths possible.

(5) We decide acceptance of composed paths as in Definition 2.3 (page 15) , where preclusion uses accepted paths for deciding.

Note that we reason in this system not only with, but also about relative strength of truth values, which are just nodes, this is then, of course, used in the acceptance condition, in preclusion more precisely.

5.4 Inheritance as reasoning with prototypes

LABEL: Section 4.4

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 “*not 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. 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 2.3 (page 19) . 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 (page 11) . 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.

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, i.e. possibly with exceptions) inclusion in a set or its complement.

In this reading, specificity stands for soft, i.e. possibly with exceptions, set inclusion. So, if b and c are visible from a , and there is a valid path from c to b , 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. This results in a partial, and not necessarily transitive relation between these distances.

When we go now from b to c , we lose 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.

We make now 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 authors are 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 a remark on prototypes.

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.

karl-search= End Is-Interpr

6 Detailed translation of inheritance to modified systems of small sets

karl-search= Start Is-Trans

LABEL: Section Translation LABEL: Section Is-Trans

6.1 Normality

LABEL: Section 5.1

As we saw already in Section 5.2 (page 35) , normality in inheritance (and Reiter defaults etc.) is relative, and as much normality as possible is preserved. There is no set of absolute normal cases of X , which we might denote $N(X)$, but only for ϕ a set $N(X, \phi)$, elements of X , which behave normally with respect to ϕ . Moreover, $N(X, \phi)$ might be defined, but not $N(X, \psi)$ for different ϕ and ψ . 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. 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. Connections between the different normalities will be established only by valid paths. 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 6.1 (page 42) .

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.

6.2 Small sets

LABEL: Section 5.2

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

As said, 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.

Consider now the system P (with Cumulativity), see Definition 4.2 (page 26) . Recall from Remark 4.1 (page 30) that small sets (see Definition 4.3 (page 29)) are used in two conceptually very distinct ways: $\alpha \vdash \beta$ 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 $\alpha \vdash \beta$ carries over to $\alpha \wedge \gamma : \alpha \wedge \gamma \vdash \beta$. (See also the discussion in [Sch04], page 86, after Definition 2.3.6.) It is this transfer which we will consider here, and not things like AND, which connect different $N(X, \phi)$ for different ϕ .

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

Fact 6.1

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

LABEL: Fact 5.1

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

Proof

(+++*** Orig.: Proof)

- (a) Consider the Tweety diagram, Diagram 2.1 (page 10) . $c \rightarrow b \rightarrow d, c \not\rightarrow d$. 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 $c \cap b < b$. $c \cap b = (c \cap b \cap d) \cup (c \cap b \cap \mathcal{C}d)$. $c \cap b \cap \mathcal{C}d \subseteq b \cap \mathcal{C}d < b$, the latter by $b \rightarrow d$, thus $c \cap b \cap \mathcal{C}d < b$, 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 $c \cap \mathcal{C}b < c$, 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 or arrows from inherited information or valid paths. In the language of big and 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 5.1 (page 38) .

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]. 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 and small subsets and their modifications is now quite complicated: we seem to have to relativize, and we seem to need two types of big and 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 and 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 $X \cap Y$.

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

(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 would result in 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.

(Here and in what follows, we are 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.)

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 and small in various, per default independent dimensions) this is close to the reasoning with absolutely big and 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, i.e. upward chaining, we argue - using 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.

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

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

Thus:

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

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

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

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

In more complicated situations, we implement e.g. the general principle (P2.2) from Definition 2.2 (page 14) , 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 , $\sigma : A \dots \rightarrow B, B \rightarrow C$, and $\sigma' : A \dots \rightarrow B$ 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 same argument as the one for the simplified definition of preclusion - see Remark 2.1 (page 17) , (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 where we 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.

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 Definition 2.3 (page 15) .

Fact 6.2

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

LABEL: Fact 5.2

The above definition and the one outlined in Definition 2.3 (page 15) correspond.

Proof

(+++*** Orig.: 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 $X \cap CZ$.

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

The first hypothesis of (2'), $X \cap \bigcap Y_i \cap CZ \in b(X \cap \bigcap Y_i, Z)$ 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 $\bigcap Y_i \cap CZ \in B(\bigcap Y_i, Z)$ must have been deduced, which can only be by (1.3'), but the set of Y_i there was chosen to take all Y_i into account for which there is a valid path from X to Y_i and arrows from the Y_i to Z (the rule was only present for the most specific reference class with respect to X and $Z!$), and we are done by the definition of valid paths in Section 2 (page 9) .

□

We summarize our ingredients.

Inheritance was done essentially by (1) and (2) of (A) above 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 and 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, where the arrows code dependencies between the situations for the evaluation of the modal quantifiers. In our picture, information at A can become

relevant only to node B iff there is a valid positive path from B to A . 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 in our translation to small and big sets formally by the possibility of combining different small and big sets in information transfer. This is, of course, a special form of relevance, there might be other forms.

karl-search= End Is-Trans

7 Conclusion

We presented a particularly simple inheritance algorithm, interpreted it in terms of information and information transfer, and the diagram itself as truth values (points) and their comparisons (through valid positive paths).

We compared our approach to the ideal reasoning with small and big sets of preferential structures.

We also showed the usefulness of applying reactive diagrams to inheritance nets: they can code the algorithm, and memorize results for future use, simplify and visualize it.

Future research should investigate on a larger scale common points and differences of semantics based and of more procedural approaches to nonmonotonic logics.

One extension of present work is the following:

We have a concatenation algorithm, which we can generalize to nonmonotonic consequence relations: if we interpret $x \rightarrow y$ as $x \vdash y$, and $x \nrightarrow y$ as $x \not\vdash y$, then the network can represent a consequence relation. The paths become the transitive closure paths from one node to another.

- (1) It is well-known that full transitivity (+ Supraclassically) entails monotony, so we cannot have transitivity.
- (2) Our algorithm is upward chaining, so we will have to follow the direction of entailment.
- (3) The default rule is to allow concatenation (transitivity) but the algorithm controls it.
- (4) If we deduce a conflict, we might solve it:
 - If the prerequisites of the last applied rules are comparable, we let the more specific rule win. Comparison is done by the same principle: We have to be able to concatenate rules by the same principle: proving that the prerequisite of the less specific rule is a consequence of the prerequisite of the more specific rule.
 - If there is no such conflict resolution, we stop concatenating.

See [Gab08c] for more details and for connections with argumentation networks.

ACKNOWLEDGEMENTS

(+++*** Orig.: ACKNOWLEDGEMENTS)

LABEL: Section 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 a former PhD student J.Ben-Naim in multi-valued and information source logic has certainly influenced us 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
- [Ant05] A.Antonelli, “Grounded consequence for defeasible reasoning”, Cambridge University Press, 2005
- [Ant97] A.Antonelli, “Defeasible inheritance on cyclic networks”, *Artificial Intelligence*, vol. 92 (1997), pp. 1-23
- [Ant99] A.Antonelli, “A directly cautious theory of defeasible consequence for default logic via the notion of general extensions”, *Artificial Intelligence*, vol. 109 (1999), pp. 71-109
- [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
- [Gab04] D.M.Gabbay, “Reactive Kripke semantics and arc accessibility”, *Proceedings CombLog04*, W.Carnielli, F.M.Dionesio, P.Mateus Eds., Centre of Logic and Computation, University of Lisbon 2004, pp. 7-20
- [Gab08b] D.M.Gabbay, “Reactive Kripke semantics and arc accessibility”, In *Pillars of Computer Science: Essays dedicated to Boris (Boaz) Trakhtenbrot on the occasion of his 85th birthday*, A.Avron, N.Dershowitz, A.Rabinovich eds., LNCS, vol. 4800, Springer, Berlin, 2008, pp. 292-341
- [Gab08c] D.M.Gabbay, “Logical modes of attack in argumentation networks”, in preparation
- [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