The Cause and Treatments of Floating Conclusions and Zombie Paths*

Yi Mao[1,2], Beihai Zhou[1,3]
[1] Institute of Logic and Cognition, Sun Yat-sen University, China
[2] Smart Card Research, Gemalto Inc.
8311 North FM 620, Austin, TX 78726, U.S.A
yi.mao@gemalto.com
[3] Philosophy Department, Peking University
Beijing, 100871, China
zhoubh@phil.pku.edu.cn

Abstract

This paper analyzes the cause of floating conclusions and zombie paths that the "directly skeptical" approach to defeasible inheritance nets encounters. Two treatments are proposed to resolve these two deep difficulties. One treatment minimally revises the original definition of a permitted path that Horty *et al.* have given. The other is based on an in-depth understanding of the basic features of nonmonotonic reasoning and the idea that nonmonotonic reasoning consists of a first stage to deduce all default conclusions and a second stage to select global conclusions from default ones by eliminating conflicts. The second treatment is much simpler and more flexible with principles that resolve conflicts. It is concluded that the differences between nonmonotonic reasoning and classical deduction are more fundamental than their apparent similarities. Horty *et al.* have pushed the analogy between two types of reasoning too far.

1 Introduction

1.1 Motivation

In the research of nonmonotonic reasoning, there has been a tendency to draw analogies between nonmonotonic reasoning and monotonic reasoning. The analogies between nonmonotonic reasoning and monotonic reasoning have been pushed to an extreme in the theory of inheritance nets with forwarding chains ([Horty et al.,1990]). Hereafter, we refer it as Horty's theory for short. This theory compares a path in a net as a proof in a classic logic system. A permitted path is defined inductively in terms of its initial segments being permitted together with some other conditions, in a similar manner as defining a proof.

However, nonmonotonic reasoning, as its name indicates, is different from monotonic reasoning. The former must take all premises into account to infer a conclusion, because any additional premise may defeat the conclusion that could be otherwise drawn. The latter does not have to fulfill this requirement. In monotonic reasoning, the conclusion from any part of the premises remains a conclusion of all premises. The gap between the differences in two types of reasoning and the analogous treatments results in unwelcome consequences of the inductive approach that Horty *et al.* take.

Makinson and Schlechta [1991] announced that Horty's theory has two deep and insuperable obstacles: floating conclusions and zombie paths. They pointed two causes for these two difficulties, and claimed that they cannot be overcome inside the theory. We agree with the criticism that Makinson and Schlechta raised against Horty's theory. However, we intend to further reveal that two difficulties share the same root cause – the inductive approach that is driven by the analogies between two types of reasoning.

We will provide two solutions to remove these difficulties. One minimally revises the original definition of a permitted path that Horty *et al.* have given. This solution is specific to Horty's theory, and is only aimed to negatively answer Makinson and Schlechta's prediction that the problems cannot be fixed inside the framework of Horty's theory. The other is based on an in-depth understanding of the basic features of nonmonotonic reasoning and completely abandons the inductive approach. The second solution is a general one, and applies to other theories like argumentation systems ([Dung, 1995]) that share the similar notions to those in inheritance nets.

* This research is supported by the Project of Key Research Institute of Humanities and Social Sciences in Universities, Ministry of Education in China, grant No. 02JAZJD720017.

Through analyzing the root cause of difficulties that Horty's theory faces, the important message we intend to deliver is that the differences between two types of reasoning are more fundamental than their similarities. The analogy between two should not go too far. This view we believe has significance in guiding how to formalize non-monotonic reasoning.

To facilitate the analysis we try to pursue, some notations and basic concepts used in the inheritance nets are introduced in the next section. However, this paper is not intended to be self-contained. Certain familiarity with at least ([Horty et al.,1990] is assumed.

1.2 Notation and Basic Concepts

Following the notation used in [Horty et al., 1990], lower case English letters like a, b, …, x, y, are used for nodes; lower case Greek letters like σ and τ are used for sequences of links that connect nodes; and capital Greek letters like Γ and Δ denote nets, which are sets of nodes and links.

A direct link in a net has the form $x \to y$ or $x \nrightarrow y$, where x is an object or a kind, and y is a kind. Direct links express assertions. When x is an object, $x \to y$ is analogous to an atomic statement like "Tweety is a bird." When x is a kind, $x \to y$ represents a generic statement such as "Birds fly." $x \to y$ and $x \nrightarrow y$ consist of a pair of conflicting assertions, which does not have to be a logically contradictory pair like "Tweety is a bird" and "Tweety is not a bird" but could be conflicting generics such as "Birds fly" and "Birds do not fly."[1]

Paths are defined inductively as follows: each direct link is a path; and if $\sigma \to p$ is a path, then both $\sigma \to p \to q$ and $\sigma \to p \nrightarrow q$ are (compound) paths. A path of the form $x \to \sigma \to y$ is said to enable (or support) the assertion $x \to y$. Likewise, path $x \to \sigma \nrightarrow y$ enables $x \nrightarrow y$. Conflicting paths are those that enable a pair of conflicting assertions. An extension of a net Γ is a maximal subset of Γ that does not contain conflicting paths.

Path $x \to \tau$ is an initial segments of path $x \to \sigma \to y$ if they are identical or $x \to \tau$ is an initial segment of $x \to \sigma$. If path $x \to \tau$ is an initial segment of path $x \to \sigma \to y$, then the latter is also called a prolongation of the former.

Intuitively, direct links in the net should be permitted; compound paths assembled by adding a direct link to an already permitted path but not preempted by other paths should also be permitted. To give a formal inductive definition, paths in a net have to be linearly ordered according to some property. Horty *et al.* introduced a notion of the degree $deg_\Gamma(\sigma)$ of a path σ in net Γ. The inductive definition from [Horty et al., 1990] is repeated below. It will be referred as HTT definition hereafter.

Definition 1 (HTT definition of permission relation)
Case I: σ is a direct link. Then $\Gamma \,|\!\!> \sigma$ (read as net Γ permits path σ) iff $\sigma \in \Gamma$.
Case II: σ is a compound path with, say, $deg_\Gamma(\sigma) = n$. As an inductive hypothesis, we can suppose it is settled whether $\Gamma \,|\!\!> \sigma'$ whenever $deg_\Gamma(\sigma') < n$. There are then two sub-cases to consider, depending on the form of σ:

(1) σ is a positive path, of the form $x \to \sigma_1 \to u \to y$. Then $\Gamma \,|\!\!> \sigma$ iff
  (a) $\Gamma \,|\!\!> x \to \sigma_1 \to u$,
  (b) $u \to y \in \Gamma$,
  (c) $x \nrightarrow y \notin \Gamma$,
  (d) for all v and τ such that $\Gamma \,|\!\!> x \to \tau \to v$ with $v \nrightarrow y \in \Gamma$, there exist z, τ_1 and τ_2 such that $z \to y \in \Gamma$ and either $z = x$ or $\Gamma \,|\!\!> x \to \tau_1 \to z \to \tau_2 \to v$.

(2) σ is a negative path, of the form $x \to \sigma_1 \to u \nrightarrow y$. Then $\Gamma \,|\!\!> \sigma$ iff
  (a) $\Gamma \,|\!\!> x \to \sigma_1 \to u$,
  (b) $u \nrightarrow y \in \Gamma$,
  (c) $x \to y \notin \Gamma$,
  (d) for all v and τ such that $\Gamma \,|\!\!> x \to \tau \to v$ with $v \to y \in \Gamma$, there exist z, τ_1 and τ_2 such that $z \nrightarrow y \in \Gamma$ and either $z = x$ or $\Gamma \,|\!\!> x \to \tau_1 \to z \to \tau_2 \to v$.

The paths referred to in the inductive step always have a degree less than that of the path being considered. It is clear, in the inductive case II, that a compound path is permitted only if its initial segment has already been permitted. This requirement is too strong. We will replace it with a weaker condition in the modified definition given in Section 2.3.

Definition 2 (Acceptance of Assertion) An assertion $x \to y$ is said to be accepted by the net Γ (i.e., $\Gamma \,|\!\!> x \to y$), if there is a path $x \to \sigma \to y$ in Γ such that it enables $x \to y$ and it is permitted by Γ.

Definition 3 (Preemption): A path of the form $x \to \tau \to v \to y$ is preempted in a net Γ just in case there is a node z such that $z \nrightarrow y \in \Gamma$, and either $z = x$ or Γ permits a path of the form $x \to \tau_1 \to z \to \tau_2 \to v$. With the exact symmetry, we say that a path of the form $x \to \tau \to v \nrightarrow y$ is preempted in a net Γ just in case there is a node z such that $z \to y \in \Gamma$, and either $z = x$ or Γ permits a path of the form $x \to \tau_1 \to z \to \tau_2 \to v$.

The central intuition behind preemption is that arguments based on more specific information override arguments based on less specific information.

Definition 4 (Neutralization): A compound path is to be neutralized by any conflicting path that is not itself preempted.

The following propositions follows from Definition 1-4.

[1] In some formalisms of generics, the logical form of "Birds fly" and "Birds do not fly" are $\forall x(B(x)>F(x))$ and $\forall x(B(x)>\neg F(x))$, respectively. They are not a pair of logical contradiction.

Proposition 1 If a path $x \to \sigma \to y$ in net Γ is preempted, then it is not permitted by Γ.

Proposition 2 If a path $x \to \sigma \to y$ in net Γ is neutralized, then it is not permitted by Γ.

Proposition 3 If a path $x \to \sigma \to y$ is not permitted by net Γ, then any of its prolongations $x \to \sigma \to y \to \tau$ is not permitted either.

1.3 Two Deep Difficulties

Makinson and Schlechta ([1991]) summarized the "directly skeptical" approach as having two essential steps:

(1) a complex one: define by induction on paths, suitably ordered, a concept of "skeptically acceptable" path;

(2) a simple step: define a proposition to be acceptable iff it is accepted by some skeptically acceptable path.

It has been argued convincingly that step (2) overlooks the existence of "floating conclusions" and step (1) ignores the expected activity of "zombie paths." To set up the stage for the further discussions on these issues, we give a brief explanation of the terms "floating conclusion" and "zombie path" that are coined by Makinson and Schlechta and their formal definitions.

1.3.1 Floating Conclusions

A floating conclusion α is such a proposition that it can be reached by two conflicting and equally strong arguments. Here is an intuitive example of floating conclusion given in [Prakken, 2002]:

A-: Brygt Rykkje is Dutch since he was born in Holland.
B-: Brygt Rykkje is Norwegian since he has a Norwegian name.
A: Brygt Rykkje likes ice-skating since he is Dutch.
B: Brygt Rykkje likes ice-skating since he is Norwegain.

This example is illustrated by *Fig. 1* (arrows consisting of nets are always read upwards unless explicitly indicated to the contrary), where arguments *A* and *B* are paths $a \to f \to x \to p$ and $a \to g \to y \to p$, respectively. As *A*- and *B*- cut off each other, none of *A* and *B* is an acceptable path. However, whichever way the conflict between *A*- and *B*- is decided, we can always reach the conclusion that Brygt Rykkje likes ice-skating. Thus, $a \to p$ is a conclusion floating on arguments *A* and *B*.

Makinson and Schlechta provided a more rigorous example that satisfies the requirement of acyclic nets with only defeasible links as shown in *Fig. 2*.

In *Fig. 2*, the proposition $a \to p$ is a floating conclusion. It is supported by path $a \to f \to g \to x \to p$ in the extension that does not contain $a \to m \not\to g$. In the other extension, it is alternatively supported by path $a \to m \to n \to y \to p$. But $a \to p$ is not supported by any path common to both extensions.

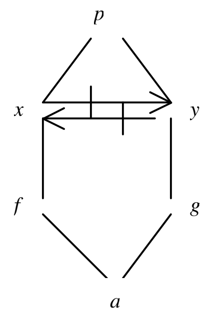

Fig. 1

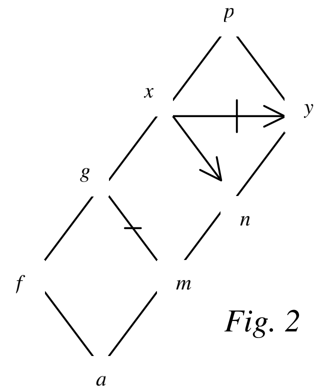

Fig. 2

Definition 5 (Floating Conclusion): A floating conclusion α is an assertion that has some supporting path in every extension, but there is no path supporting α that is common to all extensions.

HTT definition does not provide a proper treatment to floating conclusions. It does not accept a floating conclusion as shown in *Fig. 1*. It is because all supporting paths of the floating conclusion in *Fig. 1* are prolongations of some neutralized path which, according to Proposition 1, are not permitted. On the other hand, HTT definition accepts the floating conclusion "*a* is a *p*" shown in *Fig. 2* based on a wrong reason via path $a \to m \to n \to y \to p$, as Makinson and Schlechta pointed out.

When nonmonotonic reasoning is concerned, the question is whether floating conclusions should be accepted. Most scholars, like Makinson and Schlechta, think that they should be skeptically acceptable on even the most rigorous intuitive standards. It is generally regarded as a drawback of systems that do not accept floating conclusions.

Horty [2002] challenged this intuition by some intriguing counterexamples to argue that our commonly accepted inference pattern that leads to floating conclusions seems to him invalid. Prakken [2002] showed that Horty's yacht example had suppressed additional assumptions which, if made explicit, were nothing but extra information that defeats the defeasible inference. Horty's examples are not counterexamples in favor of abandoning floating conclusions, but rather these examples reflect the defeasible nature of this type of inference: conclusions may be retracted upon new information.

We second Prakken's remarks regarding nonmonotonic reasoning. It is very important to exhaustively use all known premises in a nonmonotonic reasoning when the acceptance of a conclusion is under consideration. Otherwise, a seeming counterexample is just an example showing the defeasibility of the previous conclusion drawn from fewer premises. We agree with Makinson, Schlechta and Prakken in believing that floating conclusions should be accepted, and their existence does raise a problem to "directly skeptical" approach.

1.3.2 Zombie Paths

Definition 6 Zombie paths in a net are prolongations of non-permitted paths.

Zombie paths are themselves not permitted because they contain an initial segment that is not permitted. However, they are still expected to have some negative impact on other paths. An intuitive example of zombie argument can be found in [Prakken, 2002]. To match it with the Double Diamond illustration of zombie path given in [Makinson and Schlechta, 1991], we present the example slightly modified from Prakken's:

A-: Dixon is a pacifist since he is a Quaker
B: Dixon is not a pacifist since he is a Republican
A: *A*-; and Dixon has no gun since he is a pacifist
C-: *B*; and Dixon lives in Chicago since he is a Republican
C: *C*-; and Dixon has a gun since he lives in Chicago

The graphic representation of this example is the Double Diamond shown in *Fig. 3*, where *a*, *t*, *p*, *s*, *q* and *r* stand for the following terms: Dixon, Quaker, Republican, pacifist, living in Chicago, having no gun, respectively; and arguments *A*-, *B*, *A*, *C*-, and *C* correspond to paths $a \to t \to s$, $a \to p \not\to s$, $a \to t \to s \to r$, $a \to p \to q$ and $a \to p \to q \not\to r$, respectively.

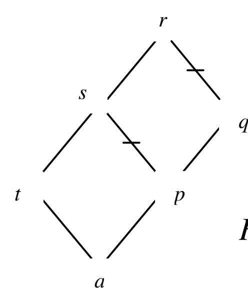

Fig. 3

In this example, the zombie path $a \to t \to s \to r$ is a prolongation of a neutralized path $a \to t \to s$. By HTT definition, we must conclude that Dixon has a gun, since the only potential counter-argument, *A*, is cut off after detecting that its sub-argument *A*- is defeated by *B*. However, *C* seems no more permissible than *A*. Though *A*- and hence *A* themselves are not permitted, they should not be completely dead for that reason alone. It is much expected that *A* should still retain sufficient power to prevent *C* from being permitted.

Theorem 1 HTT definition does not allow zombie paths to block their conflicting paths from being permitted.
Proof. Suppose that $x \to \sigma \to y$ is a zombie path in net Γ, and that it conflicts with path $x \to \tau \not\to y$. As $x \to \sigma \to y$ is a zombie path, its initial segment $x \to \sigma$ is not permitted by Γ. Under the consideration of whether $\Gamma \,|\!\!> x \to \tau \not\to y$, the zombie path $x \to \sigma \to y$ will not be taken into account by the clause Case II (2)(d) in HTT definition simply because $\Gamma \,|\!\!> x \to \sigma$ is not true. Therefore, the zombie path $x \to \sigma \to y$ does not have any impact on blocking its conflicting path $x \to \tau \not\to y$ from being permitted. The proof for the zombie paths of the form $x \to \sigma \not\to y$ symmetrically refers to the clause Case II (1)(d) in HTT definition.

Makinson and Schlechta put forth a question whether there is any way to make zombie path a genuine possibility to be able to affect adversely the strength of some other paths, within the limits of Horty's theory. We take their position in advocating that the activity of zombie paths must be well managed in any theory that deals with non-monotonic reasoning, and intend to provide a fix to the problem existing in Horty's "directly skeptical" approach.

The goals of the rest of paper are: (1) to dig out the root cause of floating conclusions and zombie paths in Horty's theory, (2) to remove that cause and fix the problems within the "directly skeptical" framework so that the change is kept minimal, (3) to discuss the nature of nonmonotonic reasoning in a general perspective and then to accordingly propose a solution to floating conclusions and zombie paths that goes beyond Horty's inductive process.

2 The Cause of Two Deep Difficulties

Makinson and Schlechta's analysis claimed that each of two steps in the architecture of "directly skeptical" approach was responsible for one of two problems respectively. However, we intend to draw attention to a common cause that triggers both floating conclusions and zombie paths.

2.1 The Analogy between Classical Deduction and Forward Chaining

Horty *et al.* have tried to push the analogy between statements that are deducible from a given set of hypotheses and assertions that are accepted by an inheritance net. This central idea is elaborated by the two-step architecture of first defining permitted paths and then defining assertions that permitted paths support. The analogy continues further on the level of permitted paths and valid arguments. The permission relation between a given net and paths that it contains is defined in an inductive manner. Mimicking the definition of a valid argument — a chain of reasoning, it is the prerequisite for constructing a compound permitted path to have permitted initial segments.

As the validity of a classical deduction is defined inductively, it is checked step by step as an argument extends so that every addition to the argument preserves validity. A full stop is called as soon as one step slides away from validity. Horty *et al.* apply the same inductive strategy to defining a permitted path. The bottom-up construction of permitted compound path interleaves with permission check on its initial segment: once a path of the form $x \to \sigma \to y$ is defeated by another path $x \to \sigma' \not\to y$ via mutual neutralization or preemption, it is fully dead. Moreover, any upwards prolongation $x \to \sigma \to y \to \tau$ of this path in the net is also simply dead. In other words, all paths that contain a defeated sub-path are completely out of the game. They do not play any role neither in supporting nor in blocking any other conclusions. Thus, some conclusions that could have been supported by upwards prolongations of some defeated paths become floating conclusions. On the other hand, some other unwanted conclusions that should have been blocked turn out to be permitted by the net.

We think that Horty *et al*. have over pushed the similarity between classical deduction and defeasible reasoning. We are in sympathy with Makinson and Schlechta's criticism to the two-step architecture resulted from the analogy. They are right to say that each of two steps in their architecture is responsible for one of two deep difficulties, respectively. However, their observation has not got all of the truth. We intend to draw attention to the root cause that brings about both problems of floating conclusions and zombie paths. The inductive method used in defining a permitted path can actually explain both difficulties that Horty's theory faces. A floating conclusion is not supported, since an argument that could have supported it is cut off at an earlier step before it reaches the floating conclusion. The inactivity of a zombie path is also due to the premature cutoff of the argument, which disables it to defeat other paths in future.

It is wrong to let the analogy run too far and carry away the objective in characterizing default reasoning. There are fundamental differences between two types of reasoning. With respect to the analogies running on two levels — upper level between statements and assertions and lower level between paths and arguments, we will present two treatments to overcome the two difficulties that Horty's theory faces. One breaks the low-level analogy between paths and deductive arguments while still keeping the upper level analogy. The other, getting closer to the characterizing features of defeasible reasoning, completely abandons the analogies altogether.

2.2 Unaccepted Assertions and Negated Defaults

In a given net, it is often the case that path $x \to \sigma \to y$ is not permitted because it is mutually neutralized by a competing path $x \to \sigma' \not\to y$. However, the death treatment of path $x \to \sigma \to y$ in Horty's theory exaggerates the fact that assertion $x \to y$ is not accepted to an extent as if it is rejected in the sense that its opposite assertion $x \not\to y$ is therefore supported, which is actually not the case.

If we interpret an assertion $x \to y$ enabled by a path of the form $x \to \sigma \to y$ in a net as a default statement, then the fact that assertion $x \to y$ is not accepted by the net may imply that its negation $\neg(x \to y)$ is consistent with the net, but it does not imply that $x \not\to y$ (i.e. $x \to \neg y$) instead is accepted. Asher and Mao [2001] had a detailed discussion regarding the proper role that a negated default should play: It should not be treated as an equivalent to $x \to \neg y$, nor should it just be ignored and does not get involved in the rest of reasoning. The power of a negated default $\neg(x \to y)$ should be just enough to block the inference of drawing the conclusion of being y because of being x but not strong enough to claim that something is not a y because it is an x. We find that this discussion on negated defaults is also applicable to the case of unaccepted assertions in an inheritance net.

In order to make clear the important difference between an unaccepted assertion and the accepted opposite assertion, we may also consider the classical provability relation as an analogy. There are three statuses between a premise set Γ and a formula α: (1) $\Gamma \vdash \alpha$; (2) $\Gamma \vdash \neg\alpha$; and (3) not ($\Gamma \vdash \alpha$) and not ($\Gamma \vdash \neg\alpha$). A premise set Γ can either prove α, or disprove it (*viz*., its negation $\neg\alpha$ is proved). But these two ends do not cover the entire spectrum. There is a middle ground where neither α nor $\neg\alpha$ is proved. The unknown status of whether α or perhaps its negation can be proved is different from the status that α is disproved. Similarly, in between the status of an assertion being accepted and being rejected by the net, there should be an intermediate unknown status, in which the assertion is neither accepted nor rejected.

As accepted assertions are defined in terms of permitted paths in Horty's theory, the inappropriate handling of unaccepted assertions is due to the inadequate treatment of non-permitted paths. Viewing non-permitted paths and hence unaccepted assertions simply as rejected ones is too coarse-grained. It is too harsh to paralyze merely non-permitted paths, while they are not actually rejected paths. Such paths are left with no power to further support any conclusions via their upwards prolongations, nor can they block any other paths that compete with their upwards prolongations. This is the common cause for both floating conclusion and zombie path problems.

A more appropriate treatment for non-permitted and non-rejected paths is to make them zombie paths: they are not "alive," but they are not fully dead either; they can still influence the status of other paths. This is what the "directly skeptical" approach needs for a repair to overcome two difficulties pointed out by Makinson and Schlechta.

Coming along with the clarification of the non-permitted (unaccepted) versus rejected paths (assertions), we propose a quick fix to HTT definition of permission relation. By weakening the requirements of having permitted paths to be non-rejected ones, the revised definition accepts floating conclusions and respects the power of zombie paths.

2.3 The Revised HTT Definition

We revise HTT definition by replacing (a) with (a′), (1)(d) with (1)(d′), and (2)(d) with (2)(d′) in Case II.

Definition 7 (Revised HTT definition)
Case I: σ is a direct link. Then $\Gamma \mid\!> \sigma$ (read as net Γ permits path σ) iff $\sigma \in \Gamma$.
Case II: σ is a compound path with, say, $deg_\Gamma(\sigma) = n$. As an inductive hypothesis, we can suppose it is settled whether $\Gamma \mid\!> \sigma'$ whenever $deg_\Gamma(\sigma') < n$. There are then two sub-cases to consider, depending on the form of σ:

(1) σ is a positive path, of the form $x \to \sigma_1 \to u \to y$. Then $\Gamma \mid\!> \sigma$ iff

(a′) for any path of the form $x \rightarrow \sigma_2 \nrightarrow u$, it is not the case that $\Gamma \mathrel{|\!\!>} x \rightarrow \sigma_2 \nrightarrow u$,

(b) $u \rightarrow y \in \Gamma$,

(c) $x \nrightarrow y \notin \Gamma$,

(d′) for all v, τ such that $x \rightarrow \tau \rightarrow v \nrightarrow y$ is a path in Γ, there exist z, τ_1 and τ_2 such that $z \rightarrow y \in \Gamma$ and either $z = x$ or $\Gamma \mathrel{|\!\!>} x \rightarrow \tau_1 \rightarrow z \rightarrow \tau_2 \rightarrow v$.

(2) σ is a negative path, of the form $x \rightarrow \sigma_1 \rightarrow u \nrightarrow y$. Then $\Gamma \mathrel{|\!\!>} \sigma$ iff

(a′) for any path of the form $x \rightarrow \sigma_2 \nrightarrow u$, it is not the case that $\Gamma \mathrel{|\!\!>} x \rightarrow \sigma_2 \nrightarrow u$,

(b) $u \nrightarrow y \in \Gamma$,

(c) $x \rightarrow y \notin \Gamma$,

(d′) for all v, τ such that $x \rightarrow \tau \rightarrow v \rightarrow y$ is a path in Γ, there exist z, τ_1 and τ_2 such that $z \nrightarrow y \in \Gamma$ and either $z = x$ or $\Gamma \mathrel{|\!\!>} x \rightarrow \tau_1 \rightarrow z \rightarrow \tau_2 \rightarrow v$.

As (a) implies (a′), condition (a′) is looser than (a). The loosened condition (a′) accommodates floating conclusions. On the other hand, since (1d′) and (2d′) imply (d) in subcases of (1) and (2) respectively, (1d′) and (2d′) are much restricted conditions that consider more competing paths before a path is to be permitted. (1d′) and (2d′) allow zombie paths to retain a power to prevent competing paths from becoming permitted. These results are stated more precisely in the following theorems.

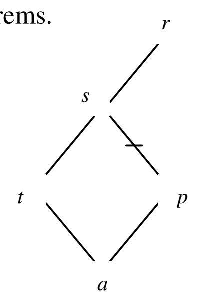

Fig. 4

Theorem 2 Proposition 2 no longer holds under the revised HTT definition.

A counterexample of Proposition 2 is illustrated by *Fig. 4*. Though path $a \rightarrow t \rightarrow s$ is neutralized by path $a \rightarrow p \nrightarrow s$, path $a \rightarrow t \rightarrow s \rightarrow r$ is permitted under Definition 7.

HTT definition guarantees that any initial segment of a permitted path is also permitted. This requirement is unnecessarily strong and fails to accept some floating conclusions and zombie paths. The revised definition ensures a weaker version of invariant: if a path is not permitted, then it must contain an initial segment that is not permitted. With the revised definition, the floating conclusion in *Fig. 1* and *Fig. 2* and zombie path in *Fig. 3* are now well taken care of.

Theorem 3 Floating conclusions are acceptable under the revised HTT definition.

As Theorem 2 supports that the prolongations of neutralized paths can now be permitted, floating conclusion shown in *Fig. 1* is accepted. Moreover, the assertion $a \rightarrow p$ in *Fig. 2* is accepted for good reason via the expected path $a \rightarrow f \rightarrow g \rightarrow x \rightarrow p$. From bottom, paths $a \rightarrow f \rightarrow g$ and $a \rightarrow m \nrightarrow g$ are mutually neutralized and hence both are not permitted. In the competition between $a \rightarrow f \rightarrow g \rightarrow x \nrightarrow y$ and $a \rightarrow m \rightarrow n \rightarrow y$, the former wins out because it preempts the latter. The conflicts between $a \rightarrow f \rightarrow g$ and $a \rightarrow m \nrightarrow g$ on node g does not have the ripple-out effect on the competing paths leading to node y that is higher up in the net. As path $a \rightarrow f \rightarrow g \rightarrow x \nrightarrow y$ is permitted, path $a \rightarrow m \rightarrow n \rightarrow y \rightarrow p$ is not permitted according to condition (a′), which is desired. On the other hand, path $a \rightarrow f \rightarrow g \rightarrow x \rightarrow p$ is permitted, because its only competing path $a \rightarrow m \rightarrow n \rightarrow y \rightarrow p$ is not considered as a genuine possibility under condition (1d′) as path $a \rightarrow f \rightarrow g \rightarrow x \nrightarrow y$ competing against its immediate segment is already supported.

Theorem 4 The revised HTT definition allows zombie paths to block their conflicting paths from being permitted.

Because the clauses Case II (1)(d′) and (2)(d′) checks to see if all of conflicting paths are preempted before the path under the consideration is to be permitted. As along as a zombie path conflicts with a path whose permissibility is examined, it will be checked and work against the conflicting path to be permitted if itself is not preempted. In *Fig. 3*, $a \rightarrow p \rightarrow q \nrightarrow r$ is not permitted under the revised HTT definition, because the zombie path $a \rightarrow t \rightarrow s \rightarrow r$ is now counted against it. The existence of path $a \rightarrow p \nrightarrow s$ does prevent path $a \rightarrow t \rightarrow s$ from being permitted. However, since $a \rightarrow p \nrightarrow s$ itself is not a permitted path either, its influence does not go beyond the node s and runs further to impact negatively on the strength of path $a \rightarrow t \rightarrow s \rightarrow r$. Looking the same problem from a different angle, though path $a \rightarrow t \rightarrow s$ is not permitted, its prolongation $a \rightarrow t \rightarrow s \rightarrow r$ nevertheless still has the right to deny a competing path $a \rightarrow p \rightarrow q \nrightarrow r$. This is exactly how a zombie path is supposed to behave.

Makinson and Schlechta have accused that the general two-step architecture of Horty's theory is just wrong. The blame is placed on the step (2): No matter how successful we might be in carrying out the complex step (1) to obtain a satisfactory definition of permitted path, floating conclusions will always stand in the way to prevent step (2) from obtaining a characterization of skeptically acceptable propositions. Based on the observation of how HTT definition accepts the floating conclusion in *Fig. 2* via an intuitively unacceptable path, they have further concluded that any attempt to accommodate floating conclusions within the limit of two-step architecture is hopeless. They have predicted that accepting floating conclusions by inflating the set of permitted paths must force one to pay the price of accepting other undesirable conclusions. However, this prediction does not apply to the revised definition that we provide.

Moreover, the revised definition is also immune to the informal proof in [Makinson and Schlechta, 1991] showing that any repair of using finitely many path values in the existing architecture cannot allow zombie paths to do their work. Our result does not mean that their proof is defective, but rather we have intentionally broken the general assumption of upwards chaining. It is not true, according to the revised definition, that a path value is always less than or equal to any of its initial segment. As a matter of fact, we allow a prolongation of a non-permitted path to be permitted, in virtue of which previous floating conclusions are now supported.

Although we fixed the HTT definition while apparently keeping two-step architecture and even the inductive structure, it may be questioned how much spirit of upwards chaining actually remains in the revised definition. The revised definition definitely does not reflect a deep-skeptical view like the original one: as soon as a counterargument is found that is at least not weaker, an argument is cut off. Prakken [2002] has made a good point that the deep skeptical approach is actually self-defeating. The earlier and more paths are cut off, the fewer paths will be defeated in a later stage, and in turn the more paths including undesired ones will have to be permitted. The intended deep skeptical approach then ends up accepting conclusions in an unexpected liberal manner. Thus, not only it does no harm to abandon deep skeptical approach, but also it helps to solve problems. It is quite convincing for us to get away from the general assumption of upwards chaining as a repair of HTT definition.

Having removed the assumption of upwards chaining, the inductive structure becomes unnecessary. It only brings in complexity without any benefits. In the next section, we will give a definition of acceptable conclusions by an inheritance net. This definition, not staying in the two-step architecture, does not depend on an inductive definition on permitted paths. The new definition lands on the nature of nonmonotonic reasoning.

3 The Basic Features of Default Reasoning

Independent of any particular approaches of formalization, Zhou and Mao [2003] have discussed some basic features regarding default reasoning based on the analysis of a set of benchmark examples. As an inheritance net is a graphic representation of default reasoning, where edges are atomic statements like "Tweety is a bird" or default statements like "Birds fly," theories of inheritance net, while characterizing default reasoning, should capture these basic features as well. We refer the reader to their paper for examples and details, but we repeat their main points here to set up a stage for proposing the second treatment to the problems of floating conclusion and zombie path.

Here is a short summary of the basic features of default (also nonmonotonic) reasoning: First of all, it is only expected to have acceptable conclusions from true premises; Secondly, the complete premise set must be taken into account to examine the acceptance of a conclusion, as it may be derivable from a subset of premises but not acceptable by the entire set of premises; Thirdly, true premises may lead to contradictory "conclusions"; Fourthly, principles like specificity are used to break contradictions; Fifthly, acceptable conclusions are deduced from the elimination of contradictory "conclusions."

While Horty's theory emphasizes the similarity between default reasoning and classical deduction, we highlight their differences. The classical deduction requires that the conclusions derived from true premises must be true. The default reasoning loosens this requirement to only having acceptable conclusions, as their conclusions may be revised or retracted in face of new information.

In classical truth-preserving deduction, if contradictions follow from a set of premises, then the premises themselves must be contradictory in the first place. On the contrary, default reasoning allows contradictions to follow from a set of non-contradictory premises like the familiar examples of Nixon Diamond and Penguin Principle. As contradictions are not acceptable but unavoidable in default reasoning, we need principles like specificity and proper methods like prioritization in default reasoning to eliminate contradictions for getting the right conclusions. Otherwise, we have to abandon default reasoning because we can by no means live with the phenomenon of "true premises but contradictory conclusions." Resolving conflicts is an important step to go through, in order to make default reasoning truly valuable for justifying some conclusions.

Due to the unique features of nonmonotonic reasoning, we think that the two-stage approach advocated in [Mao, 2003] is best suited to characterize this type of reasoning. The first stage is to boldly go ahead to get all *default conclusions* from any subset of premises. The second stage is to cautiously select *global conclusions* by eliminating contradictions. This general two-phase framework can be integrated into various approaches in the study of nonmonotonic reasoning. In authors' not yet published work, a logic system DC equipped with a possible-world semantics is constructed for the first phase to deduce all default conclusions; a preference relation among premises is used in the second phase to resolve conflicting default conclusions.

The general framework is readily applicable to inheritance nets. Every assertion that is enabled by some path is a default conclusion of the net, but it needs some work to decide which assertions can go through the contradiction elimination process and remain to be global conclusions.

4 The Definition of Acceptable Conclusions of Inheritance Nets

In the context of inheritance nets, when we say contradictions, we mean conflicting assertions of the form $x \to y$ and $x \nrightarrow y$. We are not concerned with the status of any segment of the paths that enable them. The fact that some initial segment of the path $x \to \sigma \to y$ is defeated does not necessarily ruin the acceptance of $x \to y$. A resolution is only needed at the point where two conflicting assertions compete to become the global conclusion. When there is no conflict, an assertion is naturally accepted as a global conclusion. In this picture, there is no problem to accept the floating conclusion "$a \to p$" in *Fig. 1* at all, as it even does not have a rival assertion like "$a \nrightarrow p$" to compete with. We take the strategy of resolving the conflicts at the latest moment possible so that we get the maximum set of global conclusions that can be accepted by a given net.

Definition 8 An assertion $x \to y$ is acceptable by a net Γ iff (1) there is a path $x \to \sigma \to y$ in the net; and (2) for any path such that $x \to \sigma' \to v \nrightarrow y$, there exist z, τ_1 and τ_2 such that $z \to y \in \Gamma$ and either $z = x$ or $x \to \tau_1 \to z \to \tau_2 \to v \in \Gamma$. Symmetrically, an assertion $x \nrightarrow y$ is acceptable by a net Γ iff (3) there is a path $x \to \sigma \nrightarrow y$ in the net; and (4) for any path such that $x \to \sigma' \to v \to y$, there exist z, τ_1 and τ_2 such that $z \nrightarrow y \in \Gamma$ and either $z = x$ or $x \to \tau_1 \to z \to \tau_2 \to v \in \Gamma$.

Theorem 5 Floating conclusions are acceptable under Definition 8.

Proof. If floating conclusions do not have any conflicting assertions to compete with, which is the majority case, they are straightforwardly accepted under Definition 8. In a complicated situation where a floating conclusion is challenged by a conflicting assertion, it is still acceptable if its rival is, roughly speaking[2], preempted.

Theorem 6 Zombie paths[3] can block their conflicting paths from being permitted under Definition 8.

Proof. Definition 8 examines any conflicting assertions for the acceptability of a given assertion, regardless whether it is supported by a zombie path. Zombie paths have equal chance to block the assertion enabled by its competing path from being accepted.

Definition 8 directly defines acceptable conclusions without going through permitted paths. It is much simpler than the revised HTT definition. Yet it accepts floating conclusions and has no problem with the zombie paths. The limitation of this definition is that specificity is used as the only criterion for resolving conflicts, which is indeed the only conflict-resolving principle used by inheritance nets.

The general two-stage framework allows to encode any conflict-resolving principles, other than specificity, that may be relevant to the domain of knowledge and the context of default reasoning (e.g. statistical, causal, moral, legal) in a partial ordering, which can be defined on the set of premises. The partial ordering represents the priority order among premises. By and large, the conclusion obtained from premises with high priority wins the battle. The limitation that Definition 8 has is not intrinsic to this two-stage approach in general. Definition 8 can be easily enhanced to adopt other principles to order the priority of premises, if there is such a need in inheritance nets.

5 Conclusions

The cause of floating conclusions and zombie paths in Horty's theory indicates that the analogy between a defeasible path and a deductive argument has been pushed too far. Two solutions are provided to solve the same set of problems. Our objective is not merely to fix HTT definition. The second solution, which completely abandons the analogies between the defeasible reasoning and classical deduction, is far more emphasized and preferred. It is a solution developed from a closer capture of the unique features of nonmonotonic reasoning. Our results show that the differences between two types of reasoning are rather more fundamental than their similarities. This view, if it has been established as we hope, can shape the various formalisms developed to characterize nonmonotonic reasoning, including not only inheritance nets but also logics of normality, argumentation systems, and many others.

References

[Asher and Mao, 2001] Nicholas Asher and Yi Mao. Negated defaults in commonsense entailment. *Bulletin of the Section of Logic*, Volume 30/1 (2001): 41-60.

[Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artificial Intelligence* 77(1995): 321-357.

[Horty *et al.*, 1990] John F. Horty, Richmond H. Thomason and David S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42 (1990): 311-348.

[Horty, 2002] John F. Horty. Skepticism and floating conclusions. *Artificial Intelligence*, 135 (2002): 55-72.

[Makinson and Schlechta, 1991] David Makinson and Karl Schlechta. Floating conclusions and zombie paths: two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets. *Artificial Intelligence*, 48 (1991): 199-209.

[Mao, 2003] Yi Mao. *A formalism for nonmonotonic reasoning encoded generics*. Ph.D. dissertation, University of Texas at Austin, 2003.

[Prakken, 2002] Henry Prakken. Intuitions and the modelling of defeasible reasoning: some case studies. *Proceedings of the Ninth International Workshop on Nonmonotonic Reasoning*, Toulouse, 2002, 91-99.

[Zhou and Mao, 2003] Beihai Zhou and Yi Mao. A base logic on default reasoning, *Philosophical Research* (*Supplemental Issue* 2003): 1-10.

[2] The word "preempted" should not be understood in a precise sense as defined in Definition 3, because the permitted path is not a defined notion here and hence preempted paths cannot be defined through them.

[3] Without the notion of permitted path being defined, zombie paths should be understood as prolongations of paths that enables unaccepted assertions.