PRINCIPIA 28(2): 229–246 (2024) doi: 10.5007/1808-1711.2024.e90479
Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

DIVERGING APPROACHES TO SKEPTICAL INFERENCE IN NON-MONOTONIC REASONING

JORGE MORALES DELGADO
Instituto de Investigaciones Filosóficas
Universidad de Costa Rica, COSTA RICA
jorge.moralesdelgado@ucr.ac.cr
https://orcid.org/0000-0001-9164-4036

Abstract. Our paper addresses the problem of a two-fold approach to skeptical inferences in the context non-monotonic logics. We tackle the problem through the analysis of ambiguous theories, such as the *Nixon Diamond*, as instantiated in non-monotonic inheritance networks, and the notion of an extension. Our paper presents a detailed description of the inner mechanisms underlying both approaches to skeptical inference, i.e. direct and indirect skepticism, and how each information processing policy is applied to ambiguous networks like the Nixon Diamond. Finally, we discuss the extent and limitation of each approach, and we propose an alternative stance towards the existence of diverging implementations to skeptical inferences in non-monotonic reasoning.

Keywords: logic • non-monotonic logic • defeasible reasoning • inheritance networks • deduction • knowledge representation

RECEIVED: 09/08/2022 REVISED: 03/06/2023 ACCEPTED: 09/12/2023

1. Introduction

Non-monotonic logics are a family of non-classical logics that shares the common feature of declining the monotony principle, and constructs various logical frameworks based on this omission (Anderson et al. 2013). The purpose of non-monotonic logics is to understand the intricate nature of defeasible reasoning. The main reason to advocate for the absence of the monotony principle lies on its inadequacy to reflect the dynamics of human commonsense reasoning (Pollock, 1995). Nevertheless, such endeavor has not been undertaken without its profound philosophical challenges.

One of the main challenges within the field emerges when different arguments provide grounds to opposite conclusions. For example, let us assume we have information about the political affiliations of a given individual, which we will refer to as Nixon. Suppose we have a source telling us that Nixon is a Quaker. From this information, we can defeasibly infer that such individual might take, in general, a pacifist stance on political matters. This represents a form of non-monotonic reasoning or a defeasible argument. Moreover, we can also have another source telling us that the same individual Nixon is a Republican. From this specific information alone, we can defeasibly infer that Nixon takes, in general, a non-pacifist stance on political matters. This second instance also represents a form of non-monotonic or defeasible argument. Nevertheless, both arguments or lines of reasoning taken together depict a scenario in which there are equally legitimate but opposing arguments leading to conflicting information. From one line of reasoning, we are to conclude that Nixon is a pacifist, from the other line of reasoning we are to conclude that he is a non-pacifist. This classic example is known as the *Nixon Diamond*, and it illustrates the problem of ambiguity in non-monotonic reasoning (Horty; Thomason; Touretzky 1990).

There are two main strategies employed to tackle this problem. First, one might be inclined to any conclusion as long as it is contained in any of the arguments involved. Following this approach, one would infer either that Nixon is a pacifist or that he is a non-pacifist. The undesired effect of this strategy is that one ends up endorsing arbitrary information (Horty, 2002). Second, one might state that extracting arbitrary conclusions is not adequate. According to a more conservative approach, one should infer only the conclusions that lie in the intersection of all the arguments or lines of reasoning involved. The former is known as a *credulous* approach, whereas the latter is referred to as a *skeptical* approach.

As it stands, one might be tempted to adopt the more conservative approach to draw defeasible inferences (Horty, 2002). Nevertheless, as its usually the case, the devil lies in the details, as there is not one but at least two different ways to implement the skeptical idea of intersecting the arguments in no-monotonic contexts (Gabbay & Schlechta 2016). Our work addresses the twofold approach to the same skeptical intuition. We tackle not only the technical differences and nuances but also the philosophical and epistemological ideas undermining each approach to skepticism within no-monotonic reasoning.

The aforesaid problem might not appear as a substantial concern, nevertheless, the literature has showed that the existence of diverging approaches to skeptical inference is at the heart of various fundamental problems in non-monotonic reasoning (Brachman & Levesque 2004). Moreover, the divergence itself has not received much attention despite its importance for many problems and debates within non-monotonic logics. As such, our work sets to address this matter by tackling the divergent nature of skeptical reasoning and more importantly undertaking the philosophical and epistemological nuances underlying such divergence.

To address the above, we will divide our paper into five main sections. First, we examine the problem of ambiguity as it occurs in non-monotonic inheritance networks. Second, we review the main approaches to ambiguous inheritance networks, namely the credulous and skeptical approaches. Third, we go over indirect skepticism's stance to the problem of ambiguity. Fourth, we examine the approach taken by direct skepticism to drawing inferences in ambiguous scenarios. Finally, we analyze the different aspects associated with each strategy and discuss the relevance of this divergence within non-monotonic inference and what this tells us about commonsense reasoning.

2. Ambiguity

To better understand the twofold approach to skepticism within non-monotonic reasoning we will analyze the problem known as ambiguity, which we have briefly mentioned previously in an informal fashion. We examine this problem through what is referred to as the Nixon Diamond, but now we express the problem in terms of non-monotonic inheritance networks, which is one the non-monotonic logical formalisms that have explored in depth this scenario (Schlechta, 1993).

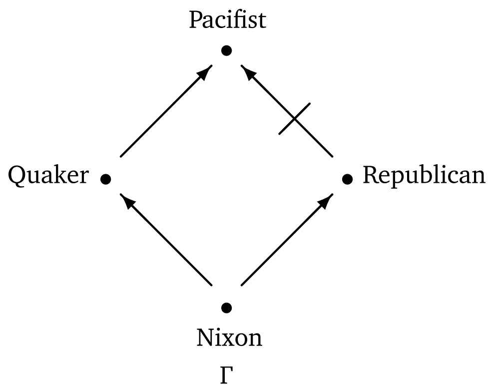

Figure 1: The Nixon Diamond

Figure 1 depicts a non-monotonic inheritance network Γ in which we have the nodes $\{$*Nixon*, *Republican*, *Quaker*, *Pacifist*$\}$ and the relations $\{$*Nixon* $\rightarrow$ *Quaker*, *Quaker* $\rightarrow$ *Pacifist*, *Nixon* $\rightarrow$ *Republican*, *Republican* $\not\rightarrow$ *Pacifist*$\}$. This specific arrangement of nodes and links in a network configuration is what is known as non-monotonic inheritance networks. The main purpose of this non-monotonic formalism is to serve as a representational framework for defeasible reasoning.

In the context of the inheritance network portraying the Nixon Diamond we can construct what we refer to as paths, which stand for defeasible arguments or lines of reasoning. In particular, we have path σ_1: *Nixon* $\rightarrow$ *Quaker* $\rightarrow$ *Pacifist* such that $\sigma_1 \in \Gamma$, and path σ_2: *Nixon* $\rightarrow$ *Republican* $\not\rightarrow$ *Pacifist* such that $\sigma_2 \in \Gamma$. We also have path σ_3: *Nixon* $\rightarrow$ *Quaker* with $\sigma_3 \in \Gamma$, as well as path σ_4: *Nixon* $\rightarrow$ *Republican* with $\sigma_4 \in \Gamma$. In a more schematic fashion, we have the following:

σ_1: *Nixon* $\rightarrow$ *Quaker* $\rightarrow$ *Pacifist*
σ_2: *Nixon* $\rightarrow$ *Republican* $\not\rightarrow$ *Pacifist*

σ_3: *Nixon* → *Quaker*
σ_4: *Nixon* → *Republican*

As we mentioned earlier, the Nixon Diamond has two important features. On the one hand, we have a path (namely σ_1) starting from the *Nixon* node that goes through the *Quaker* attribute that it is positively connected to the *pacifist* property. On the other hand, we have a path (namely σ_2) that starts with the *Nixon* node, and goes through the *Republican* attribute, which in turn, is negatively connected to the *pacifist* property. This produces a situation in which we have two equally legitimate paths pointing to conflicting information. The nature of the problem lies in the fact that the conflict surges from completely unrelated attributes, namely the *Quaker* and *Republican* attributes. Hence, with respect to the *Nixon* node, the *pacifist* property is ambiguous since both the positive and the negative connection can be drawn. Such a state of affairs is what we refer to as an ambiguity.

The aforesaid is not equivalent to the common sense of ambiguity, since everyday use might appeal to some form of vagueness or imprecision. Nevertheless, in the context of non-monotonic reasoning, an ambiguity is to be understood as a specific type of conflict, which leads to the existence of multiple extensions. In particular, an ambiguity refers to the scenario in which a starting node reaches a terminal node following two unrelated but conflicting routes, such that, *prima facie*, there is no way to resolve the conflict. The absence of a mechanism to decide how to draw information from an inheritance network that support opposing conclusions is what typifies the scenario as ambiguous.

Despite the simple structural arrangement that produces ambiguities, the problem presents a fundamental difficulty that sets the stage for opposing views concerning information processing within inheritance networks. Different stances toward this problem produce the various approaches to some of the key features of this family of knowledge representation frameworks. As Horty, Thomason & Touretzky puts it "What you say about inheritance depends crucially on your treatment of nets like the Nixon Diamond" (1990, p. 317).

3. Extensions

The existence of multiple paths within a given network raises the possibility of conflict among the various paths, i.e. in cases in which the paths support conflicting information, like the Nixon Diamond. In these scenarios there is no straightforward strategy to determine the information that we can extract. As an intermediate step to address this matter, we have at our disposal the concept of an *extension*.

An *extension* refers to a maximally consistent subset of information within a given network. That is, an extension represents a maximally non-conflicted unit of information, which is composed by *non-conflicted sets of paths* (Gabbay & Schlechta 2010). In other words, extensions represent fragments of information such that nothing else could be added without loss of consistency. The idea of an extension carries informational utility as it allows us to identify the subsets of paths that can support a given conclusion to be inferred from a network.

To better understand this notion, lets recall the Nixon Diamond example. In this non-monotonic inheritance network, we have the following paths:

$$\begin{aligned}
\sigma_1 &\ : \textit{Nixon} \to \textit{Quaker} \to \textit{Pacifist} \\
\sigma_2 &\ : \textit{Nixon} \to \textit{Republican} \not\to \textit{Pacifist} \\
\sigma_3 &\ : \textit{Nixon} \to \textit{Quaker} \\
\sigma_4 &\ : \textit{Nixon} \to \textit{Republican}
\end{aligned}$$

As we have previously stated, the problem lies on the fact that one path associated to the network provides us reasons to support that *Nixon* inherits the attribute of being a *Pacifist* (through σ_1), whereas another path within the same network justifies that *Nixon* inherits the attribute of not being a *Pacifist* (through σ_2).

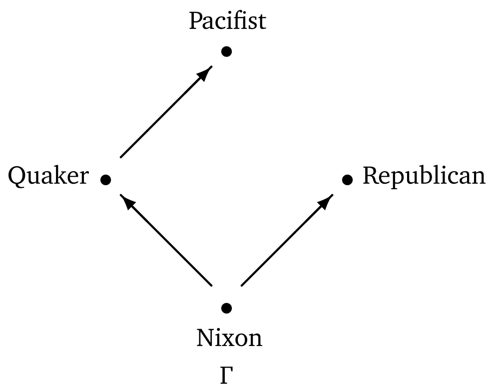

Figure 2: First Extension of the Nixon Diamond

Figure 2 portrays one of such units of information which we will refer to as Φ_1. This first extension can be characterized as follows Φ_1: {{*Nixon* → *Quaker*}, {*Nixon* → *Republican*}, {*Nixon* → *Quaker* → *Pacifist*}}, and it represents a maximally consistent unit of information associated with the network.

Figure 3 illustrates the other unit of information, which we will refer to as Φ_2. This second extension can be characterized as follows Φ_2: {{*Nixon* → *Republican*}, {*Nixon* → *Quaker*}, {*Nixon* → *Republican* $\not\to$ *Pacifist*}}. As before, Φ_2 stands as another maximally consistent unit of information associated with the network.

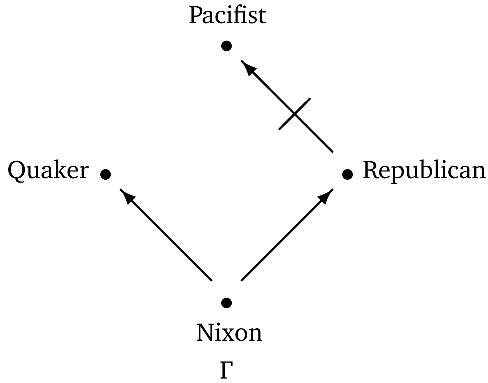

Figure 3: Second Extension of the Nixon Diamond

Based on the above, the Nixon Diamond has the following extensions:

$$\begin{aligned}\Phi_1 &= \{\{Nixon \rightarrow Quaker\},\\ &\quad \{Nixon \rightarrow Republican\},\\ &\quad \{Nixon \rightarrow Quaker \rightarrow Pacifist\}\},\end{aligned}$$

$$\begin{aligned}\Phi_2 &= \{\{Nixon \rightarrow Republican\},\\ &\quad \{Nixon \rightarrow Quaker\},\\ &\quad \{Nixon \rightarrow Republican \not\rightarrow Pacifist\}\}.\end{aligned}$$

Both sets contain every single path associated with the network such that the paths in each extension are not conflicted. The aforesaid captures the intuitive notion that when facing conflicts, we branch the information into non-conflicted units. Each extension is populated with as many non-conflicted paths as the network in question supports.

This account of an extension as an elementary unit of non-conflicted information associated with a network addresses two core tasks. On the one hand, it classifies consistent subsets of paths associated with the network. That is, an extension guarantees subsets containing information that preserves coherence and consistency. On the other hand, it captures subsets with as many paths as possible associated with the network, without losing coherence or consistency.

To summarize, an extension allows us to identify the specific subsets of paths that express the greatest amount of information associated with a network (the relative maximality property) in such a way that dismisses all the conflicting paths from the same network with the subset in question (the consistency property). As mentioned earlier, the main concern underlying the problem of ambiguity revolves around the question as to what we should infer from them. In the case of the Nixon Diamond, the question whether *Nixon* inherits the *pacifist* attribute remains open. In this sense, extensions have only captured the fragments of information within the network that are maximally consistent but are yet to solve this matter.

4. Approaches

As it turns out, there are two major strategies to address the problem of ambiguity in non-monotonic contexts based on the notion of an extension. On the one hand, we have what is referred to as a *credulous* approach. This is usually known as a more licentious strategy as it tends to infer arbitrary information. On the other hand, we have what is referred to as a *skeptical* approach. This is taken as a conservative strategy insofar as it tries to infer information, only when doing so does not incur in endorsing conflicting conclusions.

As previously mentioned, the credulous approach to inheritance structures is based on the intuition that the information that is ultimately endorsed could be any which is contained in a defeasible scenario through one of the available extensions (Nute, 2003). To address the aforementioned, the credulous approach operates on what is known as *credulous extensions*. A credulous extension is nothing more than an extension of a given network. That is, by definition *every extension of a given network is a credulous extension*. Thus, after having identified the credulous extensions of a given network, which are nothing more than the extensions of the network itself, the credulous approach states that the information that ultimately can be endorsed is the information contained in *any* of the credulous extensions (Horty 2002).

In the case of the Nixon Diamond, we have that both Φ_1 and Φ_2 are to be taken as credulous extensions of the network. Based on the credulous approach, we have equal grounds to infer that *Nixon* inherits the *pacifist* attribute or the *non-pacifist* attribute. As such, this strategy is not considered as a sound alternative and is set aside in favor of skeptical approaches (Strasser & Antonelli 2018).

The skeptical strategy is proposed as a more sensible solution to the problem of ambiguity. The underlying intuition of this approach is that the information that should be endorsed from a particular structural arrangement with multiple extensions is not whatever is contained in any of the extensions, like in the credulous stance (Brachman & Levesque 2004). To the contrary, a skeptical approach states that an attribute should be endorsed whenever such attribute is contained in the *intersection of all extensions* associated with the network. This extension-computing policy aims to provide a much more prudent scrutiny of the information it is willing to endorse. The skeptical reasoning policy stands in direct opposition to the credulous approach, as the former is much more conservative than the latter (Horty; Thomason; Touretzky 1990).

Despite the straightforward motivation behind the skeptical approach, there are two different ways to accomplish such operation. On the one hand, one could collect the attributes endorsed by all the extensions associated with a given network, *and then* apply the intersection operation on such a set of endorsed attributes. This form of skepticism is referred to as *indirect skepticism*. On the other hand, the operation of intersecting all extensions associated to the network can be applied not to the set of attributes endorsed by the extensions associated with the network, but rather to the paths contained in the extensions. This form of skepticism is referred to as *direct skepticism*. We will assess both alternatives separately.

5. Indirect Skepticism

As we previously stated, the crucial intuition underlying the skeptical approach for networks involving multiple and potentially conflicting information lies in the idea of intersecting all the extensions. The indirect version of this mechanism aims to determine when an *attribute* is common to all extensions. This approach is also referred to as an *attribute intersection* approach (Horty, 2012), but we will refer to it as the *indirect* approach (Makinson & Schlecta 1991). This approach can be divided into a three-step process.

First, we assemble all the credulous extensions associated with a given network Γ. As we already mentioned, according to credulous standards, every extension associated with a given network is a credulous extension. For such reason, skepticism (both direct and indirect) is often described as operating on credulous extensions (Gabbay & Schlechta 2010). Thus, the first step, according to direct skepticism, is defined follows:

$$\{\Phi_n : \Phi_n \text{ is an extension } \in \Gamma\}.$$

Second, and after having identified all the credulous extensions, the indirect approach proceeds to construct a further set that is composed of the attributes endorsed by every credulous extension. The aforementioned set is referred to as the set of *credulously endorsed attributes*. To get the endorsed attributes, we apply the cn() operation to the extensions associated with the network. The cn() operation is taken from Horty (2002) who employs the $*$ asterisk notation as follows:

> Where α is an argument, we will let $*\alpha$ represent the particular conclusion supported by α. Where Φ is a set of arguments, we will let $*\Phi$ represent the set of conclusions supported by the arguments in Φ—that is, the set containing the statement $*\alpha$ for each argument α belonging to Φ. (p.58)

Here, we replace the $*$ sign for a cn() operation that has the same purpose as in the context of Horty (2002), which takes either an extension or a path from a given inheritance structure and represents the information that can be derived from a given network. Thus, this second step proceeds as follows:

$$\text{cn}(\{\Phi_n \text{ is an extension } \in \Gamma\}).$$

Third, the indirect approach applies the intersection operation to the set containing the *credulously endorsed attributes*. Thus, the result of the intersection operation contains not paths (the elementary units of the extensions) but rather *attributes*. The reason why the resulting set contains attributes rather than paths lies in the fact that the set on which the intersection operation proceeds is composed of attributes rather than paths. The resulting set of attributes contains the values generated by cn() from a given network *following an indirectly skeptical approach*. The general mechanism of indirect skepticism can be summarised in the following procedure:

$$\cap(\text{cn}(\{\Phi_n \text{ is an extension } \in \Gamma\})).$$

The above states that the intersection ∩ is applied to the output of cn(). Here, the cn() operation takes paths (the base units of extensions) as inputs and returns attributes. Thus, the intersection operation computes attributes and not paths within extensions.

According to this form of skepticism the legitimacy to endorse an attribute is determined by such an attribute being common to *every* extension, i.e., such attribute is located in *every* extension. Thus, whenever an attribute is contained in every extension of a given structure, the attribute is referred to as a *skeptically endorsed statement* (Makinson & Schlecta 1991). The indirect approach ultimately seeks to identify *skeptically endorsed statements*, and only such statements are taken as skeptically acceptable conclusions.

In the context of the Nixon Diamond, the problem can be reduced to determine if either the positive or the negative attribution to the *pacifist* property is common to every credulous extension. To address the Nixon Diamond according to an indirect approach, we proceed with the same three-step process.

First, we identify all the credulous extensions, as follows:

$$\begin{aligned}\Phi_1 \quad = \quad & \{\{\textit{Nixon} \rightarrow \textit{Quaker}\},\\ & \{\textit{Nixon} \rightarrow \textit{Republican}\},\\ & \{\textit{Nixon} \rightarrow \textit{Quaker} \rightarrow \textit{Pacifist}\}\},\end{aligned}$$

$$\begin{aligned}\Phi_2 \quad = \quad & \{\{\textit{Nixon} \rightarrow \textit{Republican}\},\\ & \{\textit{Nixon} \rightarrow \textit{Quaker}\},\\ & \{\textit{Nixon} \rightarrow \textit{Republican} \not\rightarrow \textit{Pacifist}\}\}.\end{aligned}$$

Second, we construct a set containing all the attributes endorsed by each of the credulous extensions. To account for this, we make use of the cn() operation that takes extensions as input, and outputs the attributes endorsed by the extension. Thus, we apply cn() to the extensions associated with the network as follows:

$$\mathrm{cn}(\Phi_1, \Phi_2).$$

The above is just a shorter version of the following operation:

$$\begin{aligned}
&\mathrm{cn}(\{\{\textit{Nixon} \rightarrow \textit{Quaker}\},\\
&\quad \{\textit{Nixon} \rightarrow \textit{Republican}\},\\
&\quad \{\textit{Nixon} \rightarrow \textit{Quaker} \rightarrow \textit{Pacifist}\}\},\\
&\quad \{\{\textit{Nixon} \rightarrow \textit{Republican}\},\\
&\quad \{\textit{Nixon} \rightarrow \textit{Quaker}\},\\
&\quad \{\textit{Nixon} \rightarrow \textit{Republican} \not\rightarrow \textit{Pacifist}\}\})\\
&\quad = \{\{\textit{Quaker}, \textit{Republican}, \textit{Pacifist}\},\\
&\qquad \{\textit{Quaker}, \textit{Republican},\ \textit{not-Pacifist}\}\}.
\end{aligned}$$

Third, and once we have constructed the credulously endorsed attributes set, we apply the intersection operation to such set. Thus, the intersection operation is applied to the output of the cn() operation. As such, we have the following:

$$\begin{aligned}
&\cap(\{\textit{Quaker}, \textit{Republican}, \textit{Pacifist}\},\\
&\quad \{\textit{Quaker}, \textit{Republican},\ \textit{not-Pacifist}\})\\
&\quad = \{\textit{Quaker}, \textit{Republican}\}.
\end{aligned}$$

Thus, the result of the intersection operation yields the set {*Quaker*, *Republican*}, which contains nothing *with respect to the pacifist attribute*. Although the positive connection to the *pacifist* attribute is contained in the Φ_1 extension, it is not contained in the Φ_2 extension. As such, the *pacifist* attribute cannot be said to be common to all extensions. The same occurs to the negative connection of the Nixon node to the *pacifist* property, which is contained in the Φ_2 extension, but it is not contained by the Φ_1 extension. Therefore, the *pacifist* property is not common to all extensions. In such a vein, we have that neither the positive nor the negative attribution to the *pacifist* property can be inferred, and judgement on whether Nixon is a *pacifist* or not is suspended.

6. Direct Skepticism

The direct approach to skeptically process information is different from the indirect approach, in a subtle, yet significant way. Whereas the indirect approach applies the intersection operation to the set of attributes endorsed by the extensions associated with the network, direct skepticism applies the intersection operation to *paths*. Therefore, the goal from the directly skeptical point of view is to identify a *path that is common to all extensions*. Such a path is what determines the information that can be inferred from a given network. Like indirect skepticism, the direct approach can be divided into a three-step process.

First, we collect all the credulous extensions associated with a given network. As we have stated, every extension associated with a given network is, by definition, a credulous extension. Hence, this version of skepticism can also be said to operate on credulous extensions as a preliminary step. The first step according to direct skepticism is as follows:

$$\{\Phi_n : \Phi_n \text{ is an extension } \in \Gamma\}.$$

Second, once the direct account has constructed the set of all extensions associated to the network, instead of constructing a further set containing the attributes endorsed by all of the credulous extensions employing the cn() operation, it applies the intersection operation directly to the paths contained in the credulous extensions. Unlike indirect skepticism, the direct account operates on a set containing not attributes but paths (as the elementary units of extensions are paths rather than attributes). Thus, the intersection operation returns a set containing paths, not attributes. The resulting paths can be seen as the elements of what is referred to as a *skeptically endorsed extension*. This *skeptically endorsed extension* is the output of the following procedure:

$$\cap(\{\Phi_n \text{ is an extension } \in \Gamma\}).$$

Third, direct skepticism proceeds to apply the cn() operation to the aforesaid *skeptically endorsed extension*. Thus, the skeptical approach allows inferring whatever is supported by the skeptically endorsed extension. In particular, we have the following general mechanism of direct skepticism:

$$\text{cn}(\cap(\{\Phi_n \text{ is an extension } \in \Gamma\})).$$

The above states that cn() is applied to the output of the intersection operation. Like in the case of indirect skepticism, cn() acts on paths (i.e. the elementary components of extensions) rather than attributes. Nevertheless, in the case of direct skepticism it operates on what is referred to as a *skeptically endorsed extension*. This subtle distinction is what sets apart both approaches to skeptical inference.

Following direct skepticism, the ambiguity instantiated by the Nixon Diamond is reduced to determining if either the positive or the negative attribution to the *Pacifist* property is endorsed by a path common to all credulous extensions of the network. To address this inheritance network, according to a directly skeptical stance, we follow the same three-step process.

First, we identify all the credulous extensions. In the context of the Nixon Diamond, we have the following extensions:

$$\begin{aligned}\Phi_1 \quad &= \quad \{\{\mathit{Nixon} \rightarrow \mathit{Quaker}\},\\ &\qquad \{\mathit{Nixon} \rightarrow \mathit{Republican}\},\\ &\qquad \{\mathit{Nixon} \rightarrow \mathit{Quaker} \rightarrow \mathit{Pacifist}\}\},\\ \Phi_2 \quad &= \quad \{\{\mathit{Nixon} \rightarrow \mathit{Republican}\},\\ &\qquad \{\mathit{Nixon} \rightarrow \mathit{Quaker}\},\\ &\qquad \{\mathit{Nixon} \rightarrow \mathit{Republican} \not\rightarrow \mathit{Pacifist}\}\}.\end{aligned}$$

Second, and unlike the indirect approach, which extracts the credulously endorsed statements from the credulous extensions, the direct approach intersects the set containing the credulous extensions, to determine the *skeptically endorsed extension*, i.e., the set consisting of the following:

$$\begin{aligned}&\{\{\ \mathit{Nixon} \rightarrow \mathit{Quaker} \rightarrow \mathit{Pacifist}\},\\ &\ \{\ \mathit{Nixon} \rightarrow \mathit{Republican} \not\rightarrow \mathit{Pacifist}\}\}.\end{aligned}$$

Thus, the operation would proceed as follows:

$$\cap(\Phi_1, \Phi_2) = \{\{\mathit{Nixon} \rightarrow \mathit{Republican}\}, \{\mathit{Nixon} \rightarrow \mathit{Quaker}\}\}.$$

The above is just a shorter version of the following operation:

$$\begin{aligned}&\cap(\{\{\mathit{Nixon} \rightarrow \mathit{Quaker}\},\\ &\quad \{\mathit{Nixon} \rightarrow \mathit{Republican}\},\\ &\quad \{\mathit{Nixon} \rightarrow \mathit{Quaker} \rightarrow \mathit{Pacifist}\}\},\\ &\quad \{\{\mathit{Nixon} \rightarrow \mathit{Republican}\},\\ &\quad \{\mathit{Nixon} \rightarrow \mathit{Quaker}\},\\ &\quad \{\mathit{Nixon} \rightarrow \mathit{Republican} \not\rightarrow \mathit{Pacifist}\}\})\\ &\quad = \{\{\mathit{Nixon} \rightarrow \mathit{Republican}\}, \{\mathit{Nixon} \rightarrow \mathit{Quaker}\}\}.\end{aligned}$$

The process of intersecting the credulous extensions involved in the Nixon Diamond returns to the $\{\{\mathit{Nixon} \rightarrow \mathit{Republican}\}, \{\mathit{Nixon} \rightarrow \mathit{Quaker}\}\}$ set. This result lies in the fact that only paths σ_3 and σ_4 are paths common to all *extensions*.

Finally, we apply the cn() operation to the output of the $\cap$ operation. Unlike the indirect approach, what we take as input is the result of the intersection of the credulous extensions rather than every credulous extension. To put it another way, the cn() now operates on the *skeptically endorsed extension* rather than on the collection of credulous extensions. Thus, we have the following:

$$\begin{aligned}&\text{cn}(\{\mathit{Nixon} \rightarrow \mathit{Republican}\},\\&\quad\{\mathit{Nixon} \rightarrow \mathit{Quaker}\})\\&\quad = \{\mathit{Republican}, \mathit{Quaker}\}.\end{aligned}$$

The previous operation yields the $\{\mathit{Republican}, \mathit{Quaker}\}$ set as a result. Like the indirect approach, we first identify the credulous extensions of the network, i.e., both Φ_1 and Φ_2 are credulous extensions of the *Nixon Diamond*. Second, and unlike indirect skepticism, which further extracts the credulously endorsed statements from the credulous extensions, the direct approach immediately intersects the set containing the credulous extensions to determine the skeptically endorsed extension. In the context of the Nixon Diamond, only σ_3 and σ_4 are paths that are common to all extensions.

Based on the above, and from a directly skeptical stance, the result of intersection of the extensions in the Nixon Diamond is the $\{\{\mathit{Nixon} \rightarrow \mathit{Republican}\}, \{\mathit{Nixon} \rightarrow \mathit{Quaker}\}\}$ set, which in turn would be our skeptically endorsed extension. Consequently, the cn() operates on the aforesaid set and yields $\{\mathit{Republican}, \mathit{Quaker}\}$ as the resulting set. Thus, from a directly skeptical account, we have that, *with respect to the pacifist attribute*, the *Nixon* node inherits neither the positive connection nor the negative connection to the *pacifist* attribute. The reason is that, in the context of the Nixon Diamond, no path is common to all extensions, which enables a positive or negative association to such property.

7. Discussion

At this stage, we are better equipped to discuss the extent and limitations of a twofold approach to skeptical inferences in the context of no-monotonic reasoning. In particular, we want to address a preliminary remark on the inadequacy of credulous policies, but more importantly, two underlying similarities, as well as a fundamental difference between both versions of skepticism.

First of all, one of the concerns that is raised against credulous approaches lies in that the fundamental question as to how to compute information within inheritance structures remains unaddressed. In the case of the Nixon Diamond, we have seen how the *Nixon* node inherits the attribute of being a *pacifist*, and it also inherits the attribute of being a *non-pacifist*. Moreover, we have also seen that the best answer a credulous policy offers to this problem is arbitrariness (Horty, 2002). Therefore, the credulous "solution" seems to entirely bypass the problem itself, since the precise nature of the problem is deciding what should the *Nixon* node ought to inherit (either being or not being a pacifist) (Horty; Thomason; Touretzky 1990). Hence, the proposed solution of allowing either of them to be endorsed seems to elude the problem rather than addressing it as noted by Horty:

> [...] this variant of the second option also manages to sidestep our original question. We wanted to know what conclusions we should actually draw from the information provided by a default theory—whether or not, given the information from the Nixon Diamond, we should conclude that Nixon is a Pacifist, for example. But according to this variant, we are told only what there is good reason to believe—that both B(Pn) and B($\neg Pn$) are consequences of the theory, so that there is good reason to believe that Nixon is a Pacifist, but also good reason to believe that he is not. This may be useful information, but it is still some distance from telling us whether or not to conclude that Nixon is a Pacifist. (2002, p.57)

This common criticism is why this reasoning policy is not taken as a sensible solution to the problem of ambiguity. In particular, the credulous stance provides no useful decision mechanism when the inheritance structure produces a set of conflicting or mutually exclusive conclusions (Maier & Nute 2010), and leaves as the only alternative a choice that seems to be an unsound foundation to ground a generally extensible inference mechanism (Horty, 2002). Nonetheless, the adequacy of this commonly agreed upon thesis is not something we will undertake beyond this general remark, as it falls beyond the scope of this paper. Thus, when there are multiple extensions associated with a particular network, a more prudent form of handling information seems to be required if the inferences are going to have a sound justification.

In line with the concerns raised against the credulous approach, both accounts of skepticism are held as better strategies to address non-monotonic inheritance theories with multiple extensions. That is, the skeptical strategy can be seen as a refinement of the credulous approach (Schlechta, 1993). In particular, skeptical approaches (whether the direct or the indirect version) also build on top of the notion of a credulous extension. As previously mentioned, this is one of the few similarities between both versions of skepticism.

Both approaches to skeptical inference in non-monotonic contexts hold the methodological thesis that skeptically acceptable information is to be determined by the intersection operation of the involved extensions. That is, they both attempt to implement one and the same idea regarding how inheritance networks involving multiple extensions ought to be addressed. Nevertheless, this is precisely where the divergence between these two approaches emerges.

As we mentioned in previous sections, the difference between both versions of skepticism revolves as to how the idea of intersecting the credulous extensions is carried out. According to the indirect approach the input to which the intersection operation is applied to are sets containing attributes, whereas the direct approach applies the intersection operation not to attributes, but rather to paths. The result of the intersection operation is an attribute in the case of indirect skepticism, and a path in the case of direct skepticism. To put it another way, direct skepticism sanctions the endorsement of an attribute if such attribute is contained in a path that is common to all extensions (Horty, 2002)

On the one hand, the aim of direct skepticism is to identify a *skeptically endorsed extension*, also referred to as *skeptically acceptable arguments*. Then, we infer whatever is supported by such a skeptically endorsed extension (or skeptically acceptable argument). On the other, the purpose of indirect skepticism is to identify the skeptically endorsed attributes, which represent the information we are allowed to infer. The distinction between both implementations of the same skeptical idea towards information processing is subtle, as it relies on the same intuition on dealing with networks involving multiple extensions.

In cases like the Nixon Diamond, we have the same results whichever version of skeptical inference one choses, rendering them simply as different implementations to the same underlying idea. Nevertheless, there are situations in which they provide different and conflicting results. That is, there are specific structural arrangements of non-monotonic theories in which each approach to skepticism renders a different output (Makinson & Schlechta 1991). In this vein, the existence of a twofold solution to the same foundational intuition underlying skeptical approaches represents a fundamental challenge and opens a series of questions concerning different implementations of the mechanisms behind each approach. This state of affairs also raises the concern as to the extent of their divergence and what does it mean for inheritance networks as representational structures.

Nevertheless, the existence of a twofold approach to floating conclusions is an intrinsic feature associated with the skeptical stance towards information processing within non-monotonic logics. Moreover, the divergence for skeptical inferences is what renders the topic as one of vital importance, as the implementation of one or the other will produce different outcomes depending on the various characteristics of the context. In this vein, literature also shows that a consensus on the most adequate mechanisms is yet to be reached (Horty, 2012).

Having stated the above, such divergence seems more than a dispute as to which implementation of skepticism should be adopted within non-monotonic logics *simpliciter*. That is, a better approach to the two-fold mechanism to skeptical inferences should adopt the idea that each mechanism is one of various reasoning policies available within non-monotonic reasoning, and the notion of the best policy should be replaced by the concept of an adequate policy given a specific domain knowledge or epistemic circumstance.

8. Conclusions

In this paper, our analysis of skeptical reasoning has been placed in context of the problem of ambiguities within inheritance networks. That is, we have described how the existence of ambiguous situations requires non-monotonic theories to handle multiple extensions. We described the general strategies, namely, the credulous and skeptical approaches to ambiguous scenarios. The credulous strategy presents itself as more licentious, whereas skeptical approaches are regarded as a more conservative approach to such cases.

In line with the above, we focused on the existence of a twofold approach to skeptical inference in the context of no-monotonic reasoning, along with its underlying intuitions for each approach to skeptical inference. In particular, we have seen how both versions of skeptical inference processes conflicting information in cases involving ambiguities. We have also pointed out how, despite the difference in their approach to ambiguous cases, direct skepticism provides the same results as indirect skepticism for the Nixon Diamond. We have showed that this twofold approach to skeptical inferences is entrenched within non-monotonic reasoning insofar as both procedures seem to comply with the same skeptical desideratum. Moreover, the problem associated with this state of affairs lies in the existence of scenarios in which both procedures offer different results, in which case the question that this approach to multiple extensions set itself to tackle remains unaddressed.

Having stated the above, it seems quite clear that this problem is a foundational feature of non-monotonic reasoning. As such, instead of addressing this matter as a problem that needs to be solved, it is best to approach it as an intrinsic feature of non-monotonic reasoning in general. This stance towards the subject matter shifts our goals from solving the 'problem' to inquiring as to what kinds of insights this dilemma can provide concerning the nature of fallible and non-monotonic reasoning, and how either approach might be better equipped to handle certain situations depending on the specific circumstances. Thus, our work not only highlights a foundational problem within non-monotonic reasoning, but also puts forward the idea that the subject matter ought not to be taken as a problem but rather as a feature of non-monotonic reasoning. In this vein, our work seeks to expand our understanding of the divergence of skeptical inference, as a trait associated with non-monotonic reasoning itself.

Further research on the subject needs to be conducted in a way that accounts for two important factors. First, the twofold approach to skeptical inference is a given and not a hurdle, and as such it needs to be the foundation upon which further subsequent problems are to be understood. Second, the aforesaid can shed further light into other processes and properties of non-monotonic reasoning if we handle these traits of skeptical inference as an intrinsic feature rather than a problem.

In this vein, a better approach towards the discussion should involve the cognitive subtleties that lie at the heart of human common-sense reasoning, as a fundamental motivation for these kinds of non-monotonic logics in the first place. In this vein, it should be noted that the literature has taken a great deal of basic inference making situations as generating intuitions that better help us understand the problem at hand. As such, shifting our focus to the epistemic and argumentative intricacies of these non-monotonic scenarios might shed light on the technicalities that should guide the logical framework one decides to endorse, instead of forcing a logical mechanism (whether it is a credulous or skeptical one) to a broad range of non-monotonic reasoning cases.

The aforesaid nuances surrounding the diverging approaches to skeptical inferences can render better progress on the various intricacies of non-monotonic reasoning, which is the fundamental purpose of the various non-monotonic formalisms that have been constructed so far.

References

Anderson, M. L.; Gomaa, W.; Grant, J.; Perlis, D. 2013. An Approach to Human-Level Commonsense Reasoning, in K. Tanaka; F. Berto; E. Mares; F. Paoli (eds.), *Paraconsistency: Logic and Applications*, Springer Netherlands, pp.201–222. https://doi.org/10.1007/978-94-007-4438-7_12

Brachman, R.; Levesque, H. 2004. *Knowledge Representation and Reasoning* (First Edition). Morgan Kaufmann.

Gabbay, D. M.; Schlechta, K. 2010. An Analysis of Defeasible Inheritance Systems. In D. M. Gabbay; K. Schlechta (eds.), *Logical Tools for Handling Change in Agent-Based Systems* (pp.251–293). Springer. https://doi.org/10.1007/978-3-642-04407-6_9

Gabbay, D. M.; Schlechta, K. 2016. Defeasible Inheritance. In D. M. Gabbay; K. Schlechta (eds.), *A New Perspective on Nonmonotonic Logics* (pp.75–90). Springer International Publishing. https://doi.org/10.1007/978-3-319-46817-4_3

Horty, J. F. 2002. Skepticism and floating conclusions. *Artificial Intelligence*, **135**(1): 55–72. https://doi.org/10.1016/S0004-3702(01)00160-6

Horty, J. F. 2012. *Reasons as Defaults*. Oxford University Press.

Horty, J. F.; Thomason, R. H.; Touretzky, D. S. 1990. A skeptical theory of inheritance in nonmonotonic semantic networks, *Artificial Intelligence* **42**(2): 311–348. https://doi.org/10.1016/0004-3702(90)90057-7

Maier, F.; Nute, D. 2010. Well-founded semantics for defeasible logic. *Synthese*, **176**(2): 243–274. https://doi.org/10.1007/s11229-009-9492-1

Makinson, D.; Schlechta, K. 1991. Floating conclusions and zombie paths: Two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets. *Artificial Intelligence* **48**(2): 199–209. https://doi.org/10.1016/0004-3702(91)90061-N

Nute, D. 2003. Agents, Epistemic Justification, and Defeasibility. *Invited Address, 5th Augustus de Morgan Workshop*.

Prakken, H. 2002. Intuitions and the Modelling of Defeasible Reasoning: Some case studies. *Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR'2002)*. https://arxiv.org/abs/cs/0207031v1

Pollock, J. L. 1995. *Cognitive Carpentry: A Blueprint for how to Build a Person*. MIT Press.

Schlechta, K. 1993. Directly Sceptical Inheritance Cannot Capture the Intersection of Extensions. *Journal of Logic and Computation* **3**(5): 455–467.

Strasser, C.; Antonelli, G. A. 2018. Non-monotonic Logic. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2018/entries/logic-nonmonotonic/