

Defeasible Inheritance-Based Description Logics

Giovanni Casini

*Centre for Artificial Intelligence Research (CAIR)
CSIR Meraka Institute and UKZN, South Africa*

G.CASINI@CSIR.CO.ZA

Umberto Straccia

*Istituto di Scienze e Tecnologie dell'Informazione (ISTI)
CNR, Italy*

UMBERTO.STRACCIA@ISTI.CNR.IT

Abstract

Defeasible inheritance networks are a non-monotonic framework that deals with hierarchical knowledge. On the other hand, rational closure is acknowledged as a landmark of the preferential approach to non-monotonic reasoning. We will combine these two approaches and define a new non-monotonic closure operation for propositional knowledge bases that combines the advantages of both. Then we redefine such a procedure for Description Logics (DLs), a family of logics well-suited to model structured information. In both cases we will provide a simple reasoning method that is built on top of the classical entailment relation and, thus, is amenable of an implementation based on existing reasoners. Eventually, we evaluate our approach on well-known landmark test examples.

1. Introduction

The notion of *rational closure* (Lehmann & Magidor, 1992) is acknowledged as a landmark for non-monotonic reasoning due to its firm logical properties, but it has limited inference capabilities: *e.g.* an exceptional class will not inherit any of the typical properties from its superclass. Consider penguins: they are atypical birds, since they do not fly, but they still share a lot of typical properties of birds, *e.g.* they have wings. However, under rational closure we may not infer that penguins have wings. On the other hand, *Defeasible Inheritance Networks* (INs) (Horty, 1994) are a non-monotonic framework appropriate for the formalisation of hierarchical knowledge that does not have such a limitation, but exhibit questionable logical properties (see Section 3.1).

We combine these two approaches and define a new non-monotonic closure operation for propositional knowledge bases that combines the advantages of both, and then we apply the method to *Description Logics* (DLs) (Baader, Calvanese, McGuinness, Nardi, & Patel-Schneider, 2003), a family of logics that are known to be well-suited to model structured information.

1.1 Contributions and Roadmap

Section 2.1 is just a brief recap of a classical approach to inheritance nets, Horty's (1994) skeptical extension, while Section 2.2 describes the classical rational closure for propositional logic, generalising a method presented by Freund (1998). The remaining material addresses our contributions that can be summarised as follows.

1. In Section 3 we propose a new method to reason about INs that relies on the procedure for rational closure, and we present a Boolean extension of INs, called *Boolean defeasible Inheritance Networks* (BINs).
2. Using BINs, we develop in Section 4 a *defeasible inheritance-based propositional closure* that combines the advantages of both inheritance nets and rational closure.
3. Eventually, in Section 5 we apply the latter procedure to the case of *defeasible inheritance-based description logics*.

A major feature of the procedures we propose is that for propositional logic and DLs we still maintain all desired logical properties of rational closure, but with more inferential power with respect to exceptional classes. Moreover our method requires only the existence of a decision procedure for classical entailment and, thus, can be implemented on top of existing propositional SAT solvers and DL reasoners.

Please note that the present paper is a substantially revised and extended version of a previous work (Casini & Straccia, 2011). Specifically,

1. we provide an in-depth description of our reasoning model;
2. we extensively validate our approach w.r.t. to a series of landmark (test) examples illustrated by Horty (1994) and Sandewall (2010) (see Appendix A);
3. we provide computational complexity results related to our reasoning procedures; and
4. we include the proofs supporting our major claims (see Appendix B); due to the complexity of the notation, we add also a table summarising the meaning of the symbols we use more frequently (Appendix C).

1.2 Related Work

We refer to two main approaches to non-monotonic reasoning: *inheritance networks* on one hand and the *preferential approach* on the other.

Inheritance networks have been developed as a formalism for reasoning about taxonomic information. From the original ideas of Touretzky (1986), the approach has richly developed (Sandewall, 1986; Touretzky, Horty, & Thomason, 1987; Horty, Thomason, & Touretzky, 1987; Touretzky, Thomason, & Horty, 1991; Makinson, 1991; Simonet, 1996). See the works by Horty (1994) and Thomason (1992) for an overview, while Gabbay and Schlechta (2009) and Sandewall (2010) have more recently contributed to the field. In particular, in order to evaluate our proposal, we shall refer to the *skeptical* approach described by Horty (1987, 1994) as the landmark of the classical approach to inheritance networks.

On the other hand, Lehmann and Magidor’s rational closure falls into the preferential approach to non-monotonic reasoning; such an approach, since the first formulation by Shoham (1988), has become a main representative of non-monotonic reasoning (Kraus, Lehmann, & Magidor, 1990; Lehmann & Magidor, 1992; Makinson, 1994, 2005; Freund, 1998; Bochman, 2001; Rott, 2001; Schlechta, 2004), particularly appreciated for the solid logical characterization of the consequence relation.

Eventually, our proposal shall be applied in the field of description logics (Baader et al., 2003). Several non-monotonic DLs exist (Baader & Hollunder, 1993; Bonatti, Faella, &

Sauro, 2011c, 2011b; Brewka & Augustin, 1987; Britz, Heidema, & Meyer, 2008; Donini, Nardi, & Rosati, 2002; Giordano, Olivetti, Gliozzi, & Pozzato, 2009; Giordano, Gliozzi, Olivetti, & Pozzato, 2012b; Grimm & Hitzler, 2009; Knorr, Alferes, & Hitzler, 2011; Quantz & Royer, 1992; Straccia, 1993), which integrate several kinds of non-monotonic reasoning mechanism into DLs. Somewhat related to our proposal are the works by Britz et al. (2008) and Giordano et al. (2009, 2012b), as they address the application of the preferential methods into the DL framework, but they do not refer to rational closure. In a previous publication (Casini & Straccia, 2010) we present a procedure to apply rational closure to DLs, and such a procedure is at the basis of our actual proposal, but here we modify the approach in order to amplify its inferential power.

2. Preliminaries

For completeness, we start with the basic notions about INs and propositional rational closure we shall rely on.

2.1 Defeasible Inheritance Networks

In INs (Horty, 1994; Sandewall, 2010) there are classes (*nodes*), a strict subsumption relation and a defeasible subsumption relation among such classes (*links*). We shall indicate nodes by letters p, q, \dots , and we shall describe an IN by a pair $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, where \mathcal{S} is a set of *strict links*, while \mathcal{U} is a set of *defeasible links*. Every link in \mathcal{N} is a *direct* link, and it can be strict or defeasible, positive or negative. Specifically,

1. $p \Rightarrow q$: class p is subsumed by class q [positive strict link];
2. $p \not\Leftarrow q$: class p and class q are disjoint [negative strict link];
3. $p \rightarrow q$: an element of the class p is usually an element of the class q [positive defeasible link];
4. $p \not\rightarrow q$: an element of the class p is usually not an element of the class q [negative defeasible link].

Example 2.1. *The typical ‘penguin example’ could be represented as $\mathcal{N} = \langle \{p \Rightarrow b\}, \{p \not\rightarrow f, b \rightarrow f, b \rightarrow w\} \rangle$, reading b as ‘Bird’, p as ‘Penguin’, f as ‘Flying’ and w as ‘has Wings’.*

In the presence of strict links only, the subsumption relation between the classes would correspond simply to a transitive closure of the links: if p is subsumed by q and q is subsumed by r , then p is subsumed by r . Instead, the presence of defeasible links implies the possibility of potential inconsistencies in the hierarchy of classes, as in the penguin example: a transitive closure of the subsumption relation would force us to conclude at the same time that penguins are flying and non-flying creatures. Hence reasoning with inheritance networks consists mainly in deciding which conclusions have to be considered valid when faced with potential contradictions. In most of the classical approaches such decisions are based on the notions of *potential path* and *preemption* (a procedure that, given two conflicting paths, allows to choose the one resting on more specific information, invalidating the other). Among the various proposals, we shall briefly present Horty’s (1994)

classical approach, to which we shall refer as a landmark. First, we recall the notion of path that is shared by all the classical approaches to INs. We refer to the paths by means of Greek letters ($\pi, \rho, \sigma, \phi, \psi \dots$), or with tuples indicating the sequence of the nodes involved. For example, the triple $\pi = \langle p, \sigma, q \rangle$ indicates a path π that starts from the node p , passes through the path σ , and ends into the node q .

Definition 2.1 (Potential Paths, Horty, 1994, p. 117). *Given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, we can define iteratively its paths (where $--\rightarrow \in \{\rightarrow, \Rightarrow\}$ and $-/\rightarrow \in \{\nrightarrow, \not\rightarrow\}$)¹:*

- *Every direct link in \mathcal{N} is a simple path:*
 - *If $p --\rightarrow q \in \mathcal{S}$, $\langle p, q \rangle$ is a positive path.*
 - *If $p -/\rightarrow q \in \mathcal{S}$, $\langle p, q \rangle$ is a negative path.*
- *Assume a path $\pi = \langle t, \sigma, p \rangle$.*
 - *If π is positive,*
 - * *If $p --\rightarrow q \in \mathcal{S}$, $\langle \pi, q \rangle$ is a positive path.*
 - * *If $p -/\rightarrow q \in \mathcal{S}$, $\langle \pi, q \rangle$ is a negative path.*
 - *If π is negative,*
 - * *If $q \Rightarrow p \in \mathcal{S}$, $\langle \pi, q \rangle$ is a negative path.*

If a path is composed only by strict links, it is a strict path, otherwise it is a defeasible one.

A *potential path* represents a potential argumentation, and we have to decide if it is valid or not. If the path is strict, it is automatically considered as valid, otherwise, in case of potential conflicts between the conclusions of distinct defeasible paths, we have to choose which ones of them have to be considered as valid.

Using the notions of path and preemption, Horty defines an iterative construction of an *extension* of a net, that is, of the set of the paths considered valid in the net. Due to the complexity of the formal definition, we describe just roughly the procedure, referring to Horty's work (1994, sect. 3), for a better insight.

The inductive construction of the paths is based on a notion of *degree* of paths. Once we have identified the potential paths, we can define also a notion of *generalised path*. That is, a generalisation of defeasible paths that, given $\star \in \{\Rightarrow, \not\Rightarrow, \Leftarrow, \rightarrow, \nrightarrow\}$, is defined as:

1. every direct link is a generalised path;
2. if $\pi = \langle \sigma, p \rangle$ is a generalised path, and $p \star q \in \mathcal{N}$, then $\langle \pi, q \rangle$ is a generalised path.

We associate to every defeasible path $\langle p, \sigma, q \rangle$ a number that corresponds to the number of links in the longest generalised path between p and q , denoted as $deg_{\mathcal{N}}(\langle p, \sigma, q \rangle)$.

In inheritance nets, the notion of contradiction corresponds to the notion of *conflict*: we say that a path is conflicted if there is another path with the same starting and end points, but opposite polarity (*i.e.*, we have two potential paths moving from a node p to a

1. We assume the link $\not\Rightarrow$ to be symmetric, that is, if $p \not\Rightarrow q \in \mathcal{S}$, then $q \not\Rightarrow p \in \mathcal{S}$ too.

node q , but one is positive while the other is negative). If we are dealing with strict paths, the presence of a conflict points out an actual contradiction, while, dealing with defeasible paths, the contradiction is just potential and could be resolved by the notion of *preemption* (Horty, 1994, Def. 3.2.2) that allows to prefer the path relying on more specific information. In the procedure defined by Horty we start with a net \mathcal{N} , and, working iteratively on the degree of the paths, we define a sequence of pairs $\langle \mathcal{N}, \mathcal{P}_i \rangle$, where \mathcal{P}_i is the set of the paths considered valid at the i -th step. In brief, given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, Horty's procedure results into the construction of a sequence of sets of valid paths $\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2, \dots$ where:

1. $\mathcal{P}_0 = \mathcal{N}$; and
2. \mathcal{P}_{n+1} is \mathcal{P}_n united with all the paths of degree $n + 1$ and that are the extension of a valid path of degree n and that preempt all the eventual conflicting paths.

The skeptical extension of a net \mathcal{N} is defined by $\bigcup_{n=1}^{\infty} \mathcal{P}_n$ (Horty, 1994, sect. 2.2.2 and 3.3.2). To indicate that a (positive or negative, strict or defeasible) connection between two nodes p and q is considered valid in Horty's skeptical extension we write $\mathcal{N} \rightsquigarrow p \star q$, with $\star \in \{\Rightarrow, \not\Rightarrow, \rightarrow, \not\rightarrow\}$.

Example 2.2. Consider example 2.1; following Horty's procedure we obtain a skeptical extension of the net composed of the valid paths $p \Rightarrow b, p \not\rightarrow f, b \rightarrow f, b \rightarrow w, p \Rightarrow b \rightarrow w$.

Consistency. A net is considered inconsistent iff we can validate two conflicting paths, that is, there is a pair of nodes p and q such that we can derive both a positive connection ($\mathcal{N} \rightsquigarrow p \rightarrow q$ or $\mathcal{N} \rightsquigarrow p \Rightarrow q$) and a negative one ($\mathcal{N} \rightsquigarrow p \not\rightarrow q$ or $\mathcal{N} \rightsquigarrow p \not\Rightarrow q$) between them.

2.2 Propositional Rational Closure

Non-monotonic systems can be analysed from the point of view of the properties of the consequence relations they define (Makinson, 1994). From such a perspective INs do not satisfy some desirable logical properties, presented below, such as (CM) and (CT) (Makinson, 1994, pp. 56-57).

Even if the satisfaction of the structural properties we are going to present is not unanimously considered as a necessary condition for a formalization of defeasible reasoning, since most of the interesting nonmonotonic logics do not satisfy some of them, we still consider their satisfaction as a desiderata: such properties are intuitive and give back a strong logical characterization of the consequence relation, they have a solid semantic characterization based on preferential interpretation (Kraus et al., 1990; Lehmann & Magidor, 1992), and have strong connections to the classic AGM approach to belief revision (Alchourrón, Gärdenfors, & Makinson, 1985). Moreover the decision problem can be often reduced to a procedure based on the classical decision problem (as in the proposal we are going to present) allowing to implement the procedure on top of existing reasoners, and the systems based on the preferential approach rarely give back counter-intuitive results. The main problems are that the inferential power of such approaches is often too weak (we can often point out conclusions that we would like to obtain, but that the system is not able to derive), and that the preferential approach has been developed for propositional logics, and the attempts to extend it to first-order languages have turned out to be quite problematic.

The proposal we are going to present tries to overcome such inferential weaknesses, characteristic of the classical preferential approaches. After presenting it for the propositional languages, we readapt the procedure to the expressivity of DLs since, even if the preferential approach cannot be easily reformulated for first-order logics, it turns out to be still appropriate for fragments of first-order logic as DLs.

In what follows, we shall present a procedure for building the rational closure of a knowledge base using the *default-assumption approach* (Poole, 1988; Makinson, 2005); such an approach reduces the construction of the rational closure into a series of checks based on the classical consequence relations. The procedure we are presenting heavily relies on the one by Freund (1998).

Specifically, consider a classical propositional language built on a finite set of propositional letters $P = \{p_1, \dots, p_n\}$, using the classical connectives $\neg, \wedge, \vee, \supset, \equiv$; sentences will be denoted by capital letters $C, D, E \dots$, while sets of sentences by capital Greek letters Γ, Δ, \dots , and \top and \perp will have the usual meaning of “true” and “false”; in the knowledge bases, we shall indicate consequential information by means of $C \vdash D$ and $C \vdash\sim D$, respectively *strict* and *defeasible conditionals*, that have to be read respectively as ‘If C , then *always* D ’ and ‘If C , then *typically* D ’. \models denotes the classical entailment relation, and, given a set of formulae Φ or a set of strict conditionals \mathcal{T} , we indicate by \models_Φ and $\models_{\mathcal{T}}$ the monotonic entailment relations obtained adding to \models , respectively, the set of propositional formulae Φ or the set of conditionals $\{C \models D \mid C \vdash D \in \mathcal{T}\}$ as extra information; we shall use \sim also to indicate a generic non-monotonic consequence relation.

The knowledge base (KB) of an agent can be represented by means of conditionals or by means of formulae; we call *conditional knowledge base* a pair $\langle \mathcal{T}, \mathcal{D} \rangle$, where \mathcal{T} is a finite set of strict conditionals and \mathcal{D} is a finite set of defeasible conditionals.

Example 2.3. *The penguin example can be encoded as: $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ with $\mathcal{T} = \{p \vdash b\}$ and $\mathcal{D} = \{p \sim \neg f, b \sim f, b \sim w\}$.*

Another way to formalize defeasible information may be based simply on formulae, using the *default-assumption approach*: a *default-assumption knowledge base* is a pair $\langle \Phi, \Delta \rangle$, where Φ and Δ are sets of formulae representing respectively what the agent considers as necessarily true and as typically true.

Example 2.4. *The penguin example could be encoded as: $\mathcal{K} = \langle \Phi, \Delta \rangle$ with $\Phi = \{p \supset b\}$ and $\Delta = \{b \supset f, p \supset \neg f, b \supset w\}$.*

We shall use the Greek letter δ to distinguish default-assumption formulae (*i.e.*, the members of Δ). We next show how to map a conditional knowledge base into a default-assumption knowledge base (we will transform KBs of the kind of the one in Example 2.3 into KBs of the kind in Example 2.4), and then we show a simple procedure to reason within the latter, by relying only on a decision procedure for \models .

We proceed as follows: (i) we define the notions of *rational consequence relation* and *rational closure*, (ii) then, we describe a procedure to build a *rational closure* using a default-assumption knowledge base.

A consequence relation \vdash is *rational* iff it satisfies the following properties (Lehmann & Magidor, 1992):

(REF)	$C \vdash C$ for every C	Reflexivity
(CT)	$\frac{C \vdash D \quad C \wedge D \vdash F}{C \vdash F}$	Cut (Cumulative Transitivity)
(CM)	$\frac{C \vdash D \quad C \vdash F}{C \wedge D \vdash F}$	Cautious Monotony
(LLE)	$\frac{C \vdash F \quad \models C \equiv D}{D \vdash F}$	Left Logical Equivalence
(RW)	$\frac{C \vdash D \quad D \models F}{C \vdash F}$	Right Weakening
(OR)	$\frac{C \vdash F \quad D \vdash F}{C \vee D \vdash F}$	Left Disjunction
(RM)	$\frac{C \vdash F \quad C \not\vdash \neg D}{C \wedge D \vdash F}$	Rational Monotony

The first six properties, (REF)–(OR), characterise the class of the preferential consequence relations: that is, given a conditional base

$$\mathcal{D} = \{C_1 \vdash E_1, \dots, C_n \vdash E_n\},$$

we say that a conditional $C \vdash D$ is in the *preferential closure* $\mathbb{P}(\mathcal{D})$ iff it is derivable from \mathcal{D} using the rules (REF)–(OR) (Kraus et al., 1990). However, the preferential closure is generally considered inferentially too weak to be satisfactory, and so it is natural to look for stronger forms of closure.

The closure under the rule (RM) is considered, between the interesting rules, the strongest one. However, given the form of such a rule (we have a negated conditional between the premises), the rational extension of a conditional base \mathcal{D} is not unique. Indeed, we have multiple possibilities to close \mathcal{D} under such a condition: for example, if we have $\mathcal{D} = \{C \vdash F\}$, then neither $C \wedge D \vdash F$ nor $C \vdash \neg D$ are in \mathcal{D} and we can choose to add to \mathcal{D} either of them in order to satisfy (RM); moreover, as this simple example shows, it is possible that the consequence relation obtained from the intersection of different rational extensions of a knowledge base does not satisfy (RM) anymore (in this particular case, the intersection would not contain neither $C \wedge D \vdash F$ nor $C \vdash \neg D$). Hence, to define a rational closure for a conditional base \mathcal{D} , we have to choose one between the possible rational extensions of \mathcal{D} . Lehmann and Magidor have defined a rational closure operation \mathbb{R} that satisfies a set of desiderata (Lehmann & Magidor, 1992, sect. 5.1-5.3).

1. $\mathbb{P}(\mathcal{D}) \subseteq \mathbb{R}(\mathcal{D})$. That is, the conditional base \mathcal{D} and every conditional preferentially derivable from it should be in the rational closure of \mathcal{D} .
2. For every conditional of form $C \vdash \perp$, $C \vdash \perp \in \mathbb{R}(\mathcal{D})$ iff $C \vdash \perp \in \mathbb{P}(\mathcal{D})$. Analogously, for every conditional of form $\top \vdash C$, $\top \vdash C \in \mathbb{R}(\mathcal{D})$ iff $\top \vdash C \in \mathbb{P}(\mathcal{D})$. The conditionals of

form $C \sim \perp$ define what situations are simply considered impossible, while the conditionals of form $\top \sim C$ indicate what is considered typical in general. Both such kinds of information are properly managed by the preferential closure.

3. If we have that $C \sim F \in \mathbb{P}(\mathcal{D})$, $C \sim \neg D$, $C \wedge D \sim F \notin \mathbb{P}(\mathcal{D})$, we prefer a closure operation adding $C \wedge D \sim F$ instead of $C \sim \neg D$: the sense of a rule as (RM) is to employ a constrained form of monotonicity (given $C \sim F$, we add $C \wedge D \sim F$), not to arbitrarily add new defaults (the addition of $C \sim \neg D$); hence, whenever possible, given a conditional $C \sim F$ we want to consider its strengthening $C \wedge D \sim F$ instead of the unmotivated addition of a conditional $C \sim \neg D$.

We shall not describe Lehmann and Magidor's rational closure by referring to the original formulation (Lehmann & Magidor, 1992). Instead, we shall directly refer to a correspondent construction, heavily relying on the procedure defined by Freund (1998), and based on the translation of a conditional KB into a default-assumption KB, which we illustrate next.

We start with a conditional KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$. The first steps (Steps 1-3) are to define an *exceptionality ranking* for all the conditionals in the KB, following the analogous procedure by Lehmann and Magidor (1992): such a ranking will permit to distinguish correctly the strict and the defeasible knowledge contained in the KB (Step 4), since part of the strict knowledge could be implicitly contained in \mathcal{D} . This will allow us to construct the correspondent default-assumption KB (Steps 5-6). Specifically:

Step 1. We translate the strict knowledge into defeasible conditionals, that is, we move from a KB $\langle \mathcal{T}, \mathcal{D} \rangle$ to $\langle \emptyset, \mathcal{D}^* \rangle$, where

$$\mathcal{D}^* = \mathcal{D} \cup \{C \wedge \neg D \sim \perp \mid C \vdash D \in \mathcal{T}\}.$$

Intuitively, in the preferential setting, saying that $C \vdash D$ is valid is equivalent to saying that its negation is an absurdity ($\neg(C \supset D) \sim \perp$) (Bochman, 2001, sect. 6.5).

Step 2. We define $\overline{\mathcal{D}^*}$ as the set of the *materialisations* of the conditionals in \mathcal{D}^* , *i.e.*, the material implications corresponding to such conditionals:

$$\overline{\mathcal{D}^*} = \{C \supset D \mid C \sim D \in \mathcal{D}^*\}.$$

Also, we indicate by $\mathfrak{A}_{\mathcal{D}^*}$ the set of the antecedents of the conditionals in \mathcal{D}^* :

$$\mathfrak{A}_{\mathcal{D}^*} = \{C \mid C \sim D \in \mathcal{D}^*\}.$$

Step 3. We define an *exceptionality ranking* of the conditionals in \mathcal{D}^* (Lehmann & Magidor, 1992, sect. 2.6). We build such a ranking on the following notion of exceptionality.

Given a set of conditionals \mathcal{D} , a formula C is *exceptional* for \mathcal{D} iff \mathcal{D} preferentially entails $\top \sim \neg C$ (*i.e.*, $\top \sim \neg C \in \mathbb{P}(\mathcal{D})$); recall that a conditional $\top \sim C \in \mathbb{R}(\mathcal{D})$ iff $\top \sim C \in \mathbb{P}(\mathcal{D})$.

A conditional $C \sim D$ is said to be exceptional for \mathcal{D} iff its antecedent C is exceptional for \mathcal{D} . The exceptionality of a proposition can be decided based on \models only (Lehmann & Magidor, 1992, Corol. 5.22), as C is exceptional for a set of conditionals

\mathcal{D} (*i.e.*, $\top \sim \neg C \in \mathbb{P}(\mathcal{D})$) iff $\overline{\mathcal{D}} \models \neg C$, where $\overline{\mathcal{D}}$ is the set of the materialisations of the conditionals in \mathcal{D} .

Let $E(\mathfrak{A}_{\mathcal{D}})$ indicate the set of the antecedents that result exceptional w.r.t. \mathcal{D} , that is

$$E(\mathfrak{A}_{\mathcal{D}}) = \{C \in \mathfrak{A}_{\mathcal{D}} \mid \overline{\mathcal{D}} \models \neg C\},$$

and with $E(\mathcal{D})$ the exceptional conditionals in \mathcal{D} , *i.e.*,

$$E(\mathcal{D}) = \{C \sim D \in \mathcal{D} \mid C \in E(\mathfrak{A}_{\mathcal{D}})\}.$$

Obviously, for every \mathcal{D} , $E(\mathcal{D}) \subseteq \mathcal{D}$.

Step 3.1. Taking under consideration the knowledge base $\langle \emptyset, \mathcal{D}^* \rangle$, we can construct iteratively a sequence $\mathcal{E}_0, \mathcal{E}_1 \dots$ of subsets of the conditional base \mathcal{D}^* in the following way:

$$\begin{aligned} \mathcal{E}_0 &= \mathcal{D}^* \\ \mathcal{E}_{i+1} &= E(\mathcal{E}_i). \end{aligned}$$

Since \mathcal{D}^* is a finite set, the construction will terminate with an (empty or non-empty) fixed point of E , *i.e.*, a set composed only of exceptional conditionals, which materialisations negate all their own antecedents.

Step 3.2. Using such a sequence, we define a ranking function \mathfrak{r} that associates to every conditional in \mathcal{D}^* a number, representing its own level of exceptionality:

$$\mathfrak{r}(C \sim D) = \begin{cases} i & \text{if } C \sim D \in \mathcal{E}_i \text{ and } C \sim D \notin \mathcal{E}_{i+1} \\ \infty & \text{if } C \sim D \in \mathcal{E}_i \text{ for every } i. \end{cases}$$

Step 4. In Step 3 we have defined the materialisation of \mathcal{D}^* and the rank of every conditional in it. Now,

Step 4.1. we can determine if \mathcal{D}^* is inconsistent. A conditional base \mathcal{D} is inconsistent if from it we can derive the conditional $\top \sim \perp$. We know from above that a conditional of the form $\top \sim C$ is in the rational closure of \mathcal{D} iff it is in its preferential closure, that is, given the result recalled in Step 3.1, we can check the consistency of \mathcal{D}^* using $\overline{\mathcal{D}^*}$: $\top \sim \perp \in \mathbb{P}(\mathcal{D}^*)$ iff $\overline{\mathcal{D}^*} \models \perp$;

Step 4.2. if \mathcal{D}^* is consistent, we define the *background theory* $\tilde{\mathcal{T}}$ of the agent as²

$$\tilde{\mathcal{T}} = \{\top \vdash \neg C \mid C \sim D \in \mathcal{D}^* \text{ and } \mathfrak{r}(C \sim D) = \infty\}.$$

Moreover, one may verify that for every conditional in \mathcal{T} there is a logically equivalent conditional in $\tilde{\mathcal{T}}$;

2. One may easily verify the correctness of this definition referring to results in the work of Bochman (2001, sect. 7.5.3, Definition 7.5.1, the definition of *clash* on p.178, Corollary 7.5.7, Definition 7.5.2, and Lemma 7.5.5). It suffices to show that the set of the conditionals with ∞ as ranking value represents the greatest clash of \mathcal{D} (the proof is quite immediate by the definition of the exceptionality ranking).

Step 4.3. once we have $\tilde{\mathcal{T}}$, we can also identify the set of conditionals $\tilde{\mathcal{D}}$, *i.e.*, the defeasible part of the information contained in \mathcal{D}^* : *i.e.*,

$$\tilde{\mathcal{D}} = \{C \sim D \in \mathcal{D}^* \mid \mathfrak{r}(C \sim D) < \infty\} \text{ (obviously, } \tilde{\mathcal{D}} \subseteq \mathcal{D} \text{)} .$$

Essentially, so far we have moved into \mathcal{T} the non-defeasible knowledge ‘hidden’ in \mathcal{D} , obtaining a new conditional base $\langle \tilde{\mathcal{T}}, \tilde{\mathcal{D}} \rangle$. Moreover, we have the ranking values of all the conditionals in $\tilde{\mathcal{D}}$.

Step 5. Now we build the default-assumption characterization of the rational closure of $\langle \tilde{\mathcal{T}}, \tilde{\mathcal{D}} \rangle$. To do so, we translate $\tilde{\mathcal{T}}$ into a set of correspondent formulae $\tilde{\Phi}$, *i.e.*,

$$\tilde{\Phi} = \{C \mid \top \vdash C \in \tilde{\mathcal{T}}\} ,$$

and $\tilde{\mathcal{D}}$ into a sequence of default-assumptions (*i.e.*, formulae) $\tilde{\Delta}$. Specifically, given the rank value of the conditionals in $\tilde{\mathcal{D}}$, we construct a sequence of default assumptions

$$\tilde{\Delta} = \langle \delta_0, \dots, \delta_n \rangle ,$$

where n is the highest rank-value in $\tilde{\mathcal{D}}$, and

$$\delta_i = \bigwedge \{C \supset D \mid C \sim D \in \tilde{\mathcal{D}} \text{ and } \mathfrak{r}(C \sim D) \geq i\} . \quad (1)$$

Defining the default-assumptions in this way, as presented by Freund (1998), we obtain a set of default formulae, each one associated with a rank value, s.t. every default formula is classically derivable from the preceding ones, that is,

$$\delta_i \models \delta_{i+1} , \text{ for } 0 \leq i < n .$$

Step 6. Given now the background theory $\tilde{\Phi}$ and the default-assumption set $\tilde{\Delta}$, according to the steps defined so far, we associate to the agent the pair $\langle \tilde{\Phi}, \tilde{\Delta} \rangle$. Combining such steps with the main theorem in Freund’s work (1998, Thm. 24), it can be shown that the default-assumption characterisation of the agent by means of the pair $\langle \tilde{\Phi}, \tilde{\Delta} \rangle$ is equivalent to the rational closure of the pair $\langle \mathcal{T}, \mathcal{D} \rangle$ defined by Lehmann and Magidor (1992). That is,

Proposition 2.1. *Given a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$,*

$$C \sim D \in \mathbb{R}(\mathcal{K}) ,$$

where \mathbb{R} is the rational closure operation defined by Lehmann and Magidor (1992), iff

$$\{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \models D ,$$

where δ_i is the first formula in $\langle \delta_0, \dots, \delta_n \rangle$ consistent with $\{C\} \cup \tilde{\Phi}$ (we indicate it by $C \vdash_{\langle \tilde{\Phi}, \tilde{\Delta} \rangle} D$).

As a consequence, using the following knowledge base transformations

$$\langle \mathcal{T}, \mathcal{D} \rangle \rightsquigarrow \langle \emptyset, \mathcal{D}^* \rangle \rightsquigarrow \langle \tilde{\mathcal{T}}, \tilde{\mathcal{D}} \rangle \rightsquigarrow \langle \tilde{\Phi}, \tilde{\Delta} \rangle, \quad (**)$$

we can characterise the rational closure of $\langle \mathcal{T}, \mathcal{D} \rangle$ via $\langle \tilde{\Phi}, \tilde{\Delta} \rangle$ by means of Proposition 2.1, *i.e.*,

$$C \vdash D \in \mathbb{R}(\langle \mathcal{T}, \mathcal{D} \rangle) \text{ iff } C \vdash_{\langle \tilde{\Phi}, \tilde{\Delta} \rangle} D.$$

So, we have a method to decide defeasible consequence under rational closure. Specifically, given a defeasible knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$ and the propositions C and D ,

1. once for all, apply to $\langle \mathcal{T}, \mathcal{D} \rangle$ the transformations $(**)$ to obtain the defeasible knowledge base $\langle \tilde{\Phi}, \tilde{\Delta} \rangle$;
2. given C , determine δ_i as the first $(\{C\} \cup \tilde{\Phi})$ -consistent formula of the sequence $\tilde{\Delta} = \langle \delta_0, \dots, \delta_n \rangle$.
3. then decide if D follows under rational closure from C w.r.t. $\langle \mathcal{T}, \mathcal{D} \rangle$ by determining whether $\{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \models D$.

Example 2.5. Consider again the case of the penguin, with the knowledge base of Example 2.3. First (**Step1**), we move the strict knowledge in \mathcal{T} into the defeasible part, obtaining

$$\mathcal{D}^* = \{p \wedge \neg b \vdash \perp, p \vdash \neg f, b \vdash f, b \vdash w\}.$$

Then (**Step2**) we define the set of the materialisations

$$\overline{\mathcal{D}^*} = \{p \wedge \neg b \supset \perp, p \supset \neg f, b \supset f, b \supset w\},$$

and the correspondent set of antecedents

$$\mathfrak{A}_{\mathcal{D}^*} = \{p \wedge \neg b, b, p\}.$$

We use the set of materialisations $\overline{\mathcal{D}^*}$ to determine the ranking value of the formulae in $\mathfrak{A}_{\mathcal{D}^*}$ and the conditionals in \mathcal{D}^* (**Step3**), obtaining

$$\begin{aligned} 0 &= \mathfrak{r}(b) = \mathfrak{r}(b \vdash f) = \mathfrak{r}(b \vdash w) \\ 1 &= \mathfrak{r}(p) = \mathfrak{r}(p \vdash \neg f) \\ \infty &= \mathfrak{r}(p \wedge \neg b) = \mathfrak{r}(p \wedge \neg b \vdash \perp). \end{aligned}$$

So (**Step4**), we define a conditional base

$$\tilde{\mathcal{K}} = \langle \tilde{\mathcal{T}}, \tilde{\mathcal{D}} \rangle,$$

with

$$\begin{aligned} \tilde{\mathcal{T}} &= \{\top \vdash \neg p \vee b\} \\ \tilde{\mathcal{D}} &= \{p \vdash \neg f, b \vdash f, b \vdash w\} \end{aligned}$$

(since in this case the strict and the defeasible part of the conditional base were correctly separated already in the initial base \mathcal{K} , we obtain that $\tilde{\mathcal{K}}$ is the same as \mathcal{K}).

Such a conditional base is translated into a knowledge base

$$\langle \tilde{\Phi}, \tilde{\Delta} \rangle$$

(**Step5**), with

$$\begin{aligned} \tilde{\Phi} &= \{ \neg p \vee b \} \\ \tilde{\Delta} &= \{ \delta_0, \delta_1 \}, \end{aligned}$$

where

$$\begin{aligned} \delta_0 &= (p \supset \neg f) \wedge (b \supset f) \wedge (b \supset w) \\ \delta_1 &= p \supset \neg f. \end{aligned}$$

Using such default information, we conclude (**Step6**) that penguins do not fly, birds fly and birds have wings.

Remark 1. Considering Example 2.5, it would be intuitive also to conclude that penguins have wings ($p \sim w$), but in the rational closure a category that is recognized as atypical, as the category of penguins in the present case (they are birds, but they don't fly, and consequently $\tau(p) = 1$), cannot inherit any of the typical characteristics of their super classes. Hence we are not allowed to conclude that, presumably, penguins have wings. Such a weak inferential power is generally considered the main limit of the rational closure. On the other hand, as we are going to see in the next section, INs manage successfully this kind of problems.

This procedure to determine the rational closure maintains the same computational complexity as the classical decision procedure, since it is easily verified that all the transformations in (**) require at most $\mathcal{O}(|\mathcal{K}|)$ entailment tests and, given also proposition 2.1 and the fact that the strict part can encode any (monotone) propositional theory, we have that

Proposition 2.2. *Deciding propositional defeasible consequence under rational closure ($\sim_{\Phi, \Delta}$) is coNP-complete.*

Lehmann and Magidor (1992) specify also a semantic characterization of the propositional rational closure, and an alternative correspondent construction has been recently presented by Giordano et al. (2012a). If we move from propositional logic to DLs, a version of rational closure for the language \mathcal{ALC} has previously been proposed (Casini & Straccia, 2010), and such a procedure can be semantically characterized by means of preferential DL interpretations (Britz, Casini, Meyer, Moodley, & Varzinczak, 2013).

As seen above, rational closure defines a non-monotonic consequence relation with an intuitive behaviour and strong logical properties; however, as by Remark 1, it is also somewhat weak, as often there are conclusions about exceptional situations that, despite intuitive, we cannot derive. Such a behaviour is due to the fact that the procedure associates to a set of premises only those conditionals that are at least as exceptional.

Next, we are going to refine rational closure in order to avoid such a loss of inferential power w.r.t. exceptional premises. Our proposal is based on a modification of the initial knowledge base: we add new conditionals that give information about exceptional cases that

would be lost in the rational closure procedure. Such a refinement is obtained using again the ranking procedure, but applying it ‘locally’, that is, in order to decide if a conditional $C \sim D$ has to be added to our KB we apply the same procedure as in rational closure, but we consider only the information that is relevant to the inferential connection between C and D . For example, assume a knowledge base composed only of the set of conditionals $\mathcal{D} = \{p \sim q, q \sim r, q \sim t, p \sim \neg t\}$; now, following the procedure for rational closure we obtain that p is an exceptional proposition, the only one, and so we cannot derive neither $p \sim r$ nor $p \sim t$. But, while we do not want to derive $p \sim t$, as we already have $p \sim \neg t$, intuitively we do not have any reason to avoid the conclusion $p \sim r$. In fact, such a conclusion would be desirable, since p is a q , and p being an r does not generate any conflict with the rest of the information in the knowledge base.

So, our aim is to specify a way to decide which information in the KB is relevant w.r.t. a particular connection (in the above case, $p \sim r$). In order to determine such a ‘local’ relatedness we are going to consider INs: we use their ‘graphical’ characterisation in order to identify the relevant information w.r.t. the connection we want to investigate, and then we apply the ranking procedure to the pieces of information recognised as relevant.

3. Boolean Defeasible Inheritance Networks

Here we present a new decision procedures for INs, based on classical propositional decisions, that, in addition to being a main step in the nonmonotonic construction we are going to present later on, turns out to be an interesting IN decision procedure per se.

In the following, we proceed as follows. At first, we define a procedure for INs, and then we map it into propositional logic, obtaining the desired refinement of rational closure.

3.1 Exceptionality Levels in Inheritance Nets

Our first aim is to apply to INs a modified version of the decision procedure for rational closure; we do this in order to define a method for deciding validity in INs that rely on propositional calculus so to allow easily (i) to extend such a method in order to include into the language also the propositional connectives \neg, \wedge, \vee , and (ii) to integrate it with rational closure, in order to extend the inferential power of rational closure without compromising its logical properties. A non-negligible side product is that it is a propositional SAT-based reasoning procedure.

We shall briefly review the case of purely strict nets showing that this case is easily manageable using propositional calculus, and then we shall focus on mixed nets.

3.1.1 STRICT NETS

For the strict part of the nets we want to obtain the same valid connections as in all classical proposals. If a net is composed only of strict links, *i.e.*, $\mathcal{N} = \langle \mathcal{S}, \emptyset \rangle$, its valid connections and its consistency can be easily checked using propositional calculus. Indeed, define a classical propositional language ℓ using the nodes in \mathcal{N} as propositional letters (call $P_{\mathcal{N}}$ such a set of propositional letters), and \supset and \neg as connectives, and translate the set of links \mathcal{S} into the set of the corresponding propositional implications

$$\Phi = \{p \supset q \mid p \Rightarrow q \in \mathcal{S}\} \cup \{p \supset \neg q \mid p \not\Rightarrow q \in \mathcal{S}\} .$$

We indicate by l a literal in ℓ (being a *literal* a propositional letter or its negation), and we define \mathfrak{A}_Φ as the set of the antecedents of the implications in Φ , that is

$$\mathfrak{A}_\Phi = \{p \mid p \supset l \in \Phi\} .$$

Then we can derive the valid paths using Φ and classical consequence relation \models^3 .

Proposition 3.1. *Consider a net $\mathcal{N} = \langle \mathcal{S}, \emptyset \rangle$ and translate it into a set of propositional implications Φ . The following properties hold:*

1. *If \mathcal{N} is a consistent net, there is a valid strict positive (resp., negative) path $\langle p, \sigma, q \rangle$ from p to q , that is $\mathcal{N} \rightsquigarrow p \Rightarrow q$ (resp., $\mathcal{N} \rightsquigarrow p \not\Leftarrow q$), iff $\Phi \models p \supset q$ (resp., $\Phi \models p \supset \neg q$).*
2. *\mathcal{N} is inconsistent iff $\Phi \models \neg p$ for some $p \in \mathfrak{A}_\Phi$.*
3. *Deciding strict consequence can be done in polynomial time.*

So we can treat the decision problem in strict nets by means of classical propositional calculus, obtaining exactly the same valid strict paths as in the classical approaches to nets.

Note that there is a difference in the notion of inconsistency between INs and propositional logic. As seen in Section 2.1, a net is considered inconsistent if there is a node p that, simultaneously, is positively and negatively connected to another node q : p is and is not, simultaneously, a subclass of q . In the inheritance nets, such a situation is interpreted as a contradiction, while in the propositional logic the correspondent situation ($\Phi \models (p \supset q) \wedge (p \supset \neg q)$) would just force the negation of the propositional letter (*i.e.*, node) p ($\Phi \models \neg p$), that would correspond to saying that no individual can fall under the class p .

3.1.2 MIXED NETS

Now we consider nets with both strict and defeasible links. In what follows we will assume that the strict part of a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$ is inferentially closed, that is, if $\mathcal{N} \rightsquigarrow p \Rightarrow q$ (resp., if $\mathcal{N} \rightsquigarrow p \not\Leftarrow q$) then $p \Rightarrow q \in \mathcal{S}$ (resp. $p \not\Leftarrow q \in \mathcal{S}$).

Our procedure differs from the classical approaches to INs mainly because it is not based on the notion of potential path; instead, we translate the net's links into propositional formulae, and then we build an exceptionality ranking using a procedure that is similar to the one defined for rational closure. The main difference with the procedure defined for rational closure lays on the 'local' characterization of the exceptionality rankings: to check if there is a valid connection between a pair of nodes p and q we proceed in defining an exceptionality ranking of the nodes; however, we do not consider all the nodes in the net, but only those 'related' to p and q . Such a relation is determined by means of the notion of *course*, that is a generalisation of the potential path.

Roughly, courses are simply routes on the net following the direction of the arrows, without considering if each of them is a positive or a negative arrow.

Definition 3.1 (Course). *Courses are defined as follows (where $\star \in \{\Rightarrow, \not\Leftarrow, \rightarrow, \nrightarrow\}$):*

3. Note that strict links can be encoded as 2-CNF formulae, also called Krom formulae, and that the propositional 2-SAT problem is in P .

1. every link $p \star q$ in \mathcal{N} is a course $\pi = \langle p, q \rangle$ in \mathcal{N} ; and
2. if $\pi = \langle \sigma, q \rangle$ is a course and $q \star r$ is a link in \mathcal{N} that does not already appear in π , then $\pi' = \langle \pi, r \rangle$ is a course in \mathcal{N} .

The omission of repetitions in courses is needed to guarantee the finiteness of courses even if the net contains cycles. So, given a net \mathcal{N} defined by a finite number of links, there is only a finite set $C^{\mathcal{N}}$ of courses, that, in turn, are finite sequences of nodes. We denote with $C_{p,q}^{\mathcal{N}}$ the set of all the courses in \mathcal{N} going from node p to the node q , *i.e.*,

$$C_{p,q}^{\mathcal{N}} = \{ \sigma \in C^{\mathcal{N}} \mid \sigma = \langle p, \sigma', q \rangle \text{ for some } \sigma' \} .$$

We next provide a procedure that defines the validity of a defeasible connection between two nodes p and q , via a mapping into propositional logic. Given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, we define a correspondent knowledge base

$$\mathcal{K}_{\mathcal{N}} = \langle \Phi_{\mathcal{N}}, \Delta_{\mathcal{N}} \rangle ,$$

where

$$\Phi_{\mathcal{N}} = \{ p \supset q \mid p \Rightarrow q \in \mathcal{S} \} \cup \{ p \supset \neg q \mid p \not\Leftarrow q \in \mathcal{S} \}$$

and

$$\Delta_{\mathcal{N}} = \{ p \supset q \mid p \rightarrow q \in \mathcal{U} \} \cup \{ p \supset \neg q \mid p \not\Leftarrow q \in \mathcal{U} \} .$$

In the following, we may omit \mathcal{N} if it is clear from the context.

We define an ‘exceptionality ranking’ of the nodes, that depends on the decision problem with respect to p and q only.⁴

So, let

$$\begin{aligned} \Delta_{p,q} &= \{ r \supset t \mid r \rightarrow t \in \sigma, \sigma \in C_{p,q} \} \cup \\ &\quad \{ r \supset \neg t \mid r \not\Leftarrow t \in \sigma, \sigma \in C_{p,q} \} , \end{aligned}$$

and consider the set of relative antecedents (l being a literal)

$$\mathfrak{A}_{p,q} = \{ a \mid a \supset l \in \Phi_{\mathcal{N}} \cup \Delta_{p,q} \} .$$

In the following, $\models_{\Phi_{\mathcal{N}}}$ will denote the supra classical entailment relation obtained adding to \models the set of propositional formulae $\Phi_{\mathcal{N}}$ as extra axioms. For the strict part of the net, if $p \models_{\Phi_{\mathcal{N}}} q$ (resp., $p \models_{\Phi_{\mathcal{N}}} \neg q$), then we say that q (resp., $\neg q$) *follows strictly* from p in \mathcal{N} , and we indicate it by $p \vdash_{\mathcal{N}} q$ (resp., $p \vdash_{\mathcal{N}} \neg q$).

On the other hand, $\vdash_{\sim \mathcal{N}}$ will indicate the inference relation for the defeasible part, that is, $p \vdash_{\sim \mathcal{N}} q$ has to be read as ‘a member of the class p is typically also a member of the class q ’ in \mathcal{N} . Analogously for $p \vdash_{\sim \mathcal{N}} \neg q$ in the negative case.

4. This is the main difference w.r.t. the procedure for propositional rational closure: while there we rank all the information in the KB at once, here we rank only the information related to the connection we are interested in, between p and q .

Now, we use $\Phi_{\mathcal{N}}$ and $\Delta_{p,q}$ to determine the exceptionality level. If we are investigating the connection between p and q , a node in $\mathfrak{A}_{p,q}$ is exceptional if it is negated by the information contained in Φ and $\Delta_{p,q}$ (compare with Step 3.2 in Section 2.2):

$$\begin{aligned} E(\mathfrak{A}_{p,q}) &= \{a \in \mathfrak{A}_{p,q} \mid \Delta_{p,q} \models_{\Phi_{\mathcal{N}}} \neg a\} \\ E(\Delta_{p,q}) &= \{a \supset b \in \Delta_{p,q} \mid a \in E(\mathfrak{A}_{p,q})\} . \end{aligned}$$

Therefore, like Step 3.3 of Section 2.2, we build a sequence

$$\begin{aligned} \alpha_0 &= \mathfrak{A}_{p,q} \\ \alpha_i &= E(\alpha_{i-1}) , \end{aligned}$$

and the corresponding sequence

$$\begin{aligned} \mathcal{E}_0 &= \Delta_{p,q} \\ \mathcal{E}_i &= E(\mathcal{E}_{i-1}) . \end{aligned}$$

Since $\mathfrak{A}_{p,q}$ and $\Delta_{p,q}$ are finite, and for every i $\alpha_i \subseteq \alpha_{i+1}$ and $\mathcal{E}_{i+1} \subseteq \mathcal{E}_i$, the sequences terminate with an (empty or non-empty) fixed point of the function E , as in Section 2.2.

Define now a ranking function (like Step 3.4) τ that associates to every implication in $\Delta_{p,q}$ a number, representing its level of exceptionality:

$$\begin{aligned} \tau_{p,q}(a) &= i \text{ if } a \in \alpha_i \text{ and } a \notin \alpha_{i+1} \\ \tau_{p,q}(a) &= \infty \text{ if } a \in \alpha_i \text{ for all } i \\ \tau_{p,q}(a \supset b) &= i \text{ if } (a \supset b) \in \mathcal{E}_i \text{ and } (a \supset b) \notin \mathcal{E}_{i+1} \\ \tau_{p,q}(a \supset b) &= \infty \text{ if } (a \supset b) \in \mathcal{E}_i \text{ for all } i . \end{aligned}$$

Clearly, $\tau(a \supset b) = \tau(a)$ for every $a \supset b \in \Delta_{p,q}$. In the following, we assume that we do not obtain any node with a ranking value of ∞ (that is, the function E terminates with an empty set). We will see later on (Proposition 3.6) that in this latter case the net is inconsistent.

We now consider the set $\widehat{\Delta}_{p,q}$ of the implications $a \supset b \in \Delta_{p,q}$ that are at least as exceptional as p ,

$$\widehat{\Delta}_{p,q} = \{a \supset b \in \Delta_{p,q} \mid \tau(a \supset b) \geq \tau(p)\} ,$$

and eventually define

$$\begin{aligned} p \vdash_{\mathcal{N}} q &\text{ iff } \widehat{\Delta}_{p,q} \models_{\Phi_{\mathcal{N}}} p \supset q \\ p \vdash_{\mathcal{N}} \neg q &\text{ iff } \widehat{\Delta}_{p,q} \models_{\Phi_{\mathcal{N}}} p \supset \neg q . \end{aligned}$$

In the language of the nets, we indicate the inference relation generated by such a procedure by the symbol \rightsquigarrow' . That is,

$$\begin{aligned} \mathcal{N} \rightsquigarrow' p \Rightarrow q &\text{ iff } p \vdash_{\mathcal{N}} q \\ \mathcal{N} \rightsquigarrow' p \not\Leftarrow q &\text{ iff } p \vdash_{\mathcal{N}} \neg q \\ \mathcal{N} \rightsquigarrow' p \rightarrow q &\text{ iff } p \vdash_{\mathcal{N}} q \\ \mathcal{N} \rightsquigarrow' p \not\rightarrow q &\text{ iff } p \vdash_{\mathcal{N}} \neg q . \end{aligned}$$

So, given $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$ and a pair of nodes $\langle p, q \rangle$, our inference procedure for INs can be summarised as follows:

1. Close \mathcal{S} under strict validity.
2. Check if there is a direct (and hence valid) link in \mathcal{N} connecting p to q . If there is, the connection is valid. Otherwise, proceed.
3. Determine the set $C_{p,q}$ of the courses in \mathcal{N} connecting p to q , map the links in \mathcal{S} and $C_{p,q}$ into the sets of implications Φ and $\Delta_{p,q}$, define the set $\mathfrak{A}_{p,q}$ of the antecedents of the implications in $\Phi \cup \Delta_{p,q}$.
4. Determine the ranking value of every proposition in $\mathfrak{A}_{p,q}$ and every implication in $\Delta_{p,q}$.
5. Define the set $\widehat{\Delta}_{p,q}$ of the implications that are at least as exceptional as p .
6. Then decide $\mathcal{N} \rightsquigarrow' p \rightarrow q$ ($\mathcal{N} \rightsquigarrow' p \not\rightarrow q$) by determining whether $\widehat{\Delta}_{p,q} \models_{\Phi} p \supset q$ ($\widehat{\Delta}_{p,q} \models_{\Phi} p \supset \neg q$).

Please note again that we rely on a decision procedure for \models only. The examples below illustrate the behaviour of our method.

Example 3.1. Consider Example 2.1 with additional links $t \Rightarrow b$ and $t \Rightarrow p$ (read t as ‘tweety’).




Figure 1: Example 3.1

We translate the net into the following knowledge base

$$\mathcal{K} = \langle \Phi, \Delta \rangle ,$$

where

$$\Phi = \{t \supset b, t \supset p, p \supset b\}$$

and

$$\Delta = \{p \supset \neg f, b \supset f, b \supset w\} .$$

Suppose now, we want to decide if t is connected to f (i.e., Tweety flies). Since the link $b \rightarrow w$ does not appear in any course from t to f , we have

$$\begin{aligned} \Delta_{t,f} &= \{p \supset \neg f, b \supset f\} \\ \mathfrak{A}_{t,f} &= \{t, b, p\} , \end{aligned}$$

and so we obtain

$$\Delta_{t,f} \models_{\Phi} \neg p \text{ and } \Delta_{t,f} \models_{\Phi} \neg t .$$

Thus,

$$\begin{aligned} 0 &= \mathfrak{r}(b) = \mathfrak{r}(b \supset f) \\ 1 &= \mathfrak{r}(t) = \mathfrak{r}(p) = \mathfrak{r}(p \supset \neg f) \end{aligned}$$

So,

$$\widehat{\Delta}_{t,f} = \{p \supset \neg f\}$$

and, as $\widehat{\Delta}_{t,f} \models_{\Phi} t \supset \neg f$, we have

$$t \sim_{\mathcal{N}} \neg f ,$$

as expected.

As next, we ask if t is connected to w (i.e., Tweety has wings). Now, we have

$$\begin{aligned} \Delta_{t,w} &= \{b \supset w\} \\ \mathfrak{A}_{t,w} &= \{t, b, p\} . \end{aligned}$$

As $\Delta_{t,w}$ does not imply the negation of any of the members of $\mathfrak{A}_{t,w}$, we have

$$0 = \mathfrak{r}(t) = \mathfrak{r}(p) = \mathfrak{r}(b) = \mathfrak{r}(b \supset w)$$

and

$$\widehat{\Delta}_{t,w} = \Delta_{t,w} .$$

As $\widehat{\Delta}_{t,w} \models_{\Phi} t \supset w$, we have

$$t \sim_{\mathcal{N}} w ,$$

as expected.

Example 3.2. Consider the Nixon Diamond (see Figure 2), where n is ‘Nixon’, r is ‘republican’, q is ‘quaker’, and p is ‘pacifist’; it is another classical problem in nonmonotonic reasoning, that is similar to the previous one but we are not informed if a path is more specific of the other (while above the link $p \Rightarrow b$ tells us that the information about the penguins is more specific than the information about the birds). So, we do not want neither $n \rightarrow p$ nor $n \not\rightarrow p$ validated.




Figure 2: Nixon diamond.

The knowledge base \mathcal{K} corresponding to the net is composed of

$$\begin{aligned}\Phi &= \{n \supset r, n \supset q\} \\ \Delta &= \{r \supset \neg p, q \supset p\}.\end{aligned}$$

We want to check if n is connected to p . So, $\Delta_{n,p} = \Delta$, and the only negated antecedent is n ($\Delta_{n,p} \models_{\Phi} \neg n$): $\tau(q \supset p) = \tau(r \supset \neg p) = 0$ while $\tau(n) = 1$, i.e., $\hat{\Delta}_{n,p} = \emptyset$. Since $\hat{\Delta}_{n,p} \not\models_{\Phi} n \supset p$ and $\hat{\Delta}_{n,p} \not\models_{\Phi} n \supset \neg p$, we conclude that

$$\begin{aligned}n &\not\sim_{\mathcal{N}} p \\ n &\not\sim_{\mathcal{N}} \neg p.\end{aligned}$$

The two following examples illustrate that our procedure and Horty's skeptical closure, notwithstanding they often manifest similar results, do not always give back the same results, nor one is included in the other.

Example 3.3. Consider the net in Figure 3.




Figure 3: Example 3.3

We want to investigate if there is a valid connection between a and p . According to Horty's skeptical closure, we cannot conclude anything about a and p ($\mathcal{N} \not\vdash a \rightarrow p$). Instead, with our approach we obtain $a \sim_{\mathcal{K}} p$ ($\mathcal{N} \rightsquigarrow a \rightarrow p$), since we have $\tau(a) = \tau(f) = \tau(g) = \tau(x) = 1$.

Example 3.4. Consider the net in Figure 4.




Figure 4: Example 3.4

We want to investigate if there is a valid connection between p and b . According to Horty (1994) we conclude $\mathcal{N} \rightsquigarrow p \not\vdash b$, while with our approach we cannot conclude anything.

So, even if in many situations the results of the two approaches are the same, we can obtain different results with them. Such different outcomes are mainly due to a difference in how conflicts are interpreted. Consider Example 3.4, where we have an unresolved conflict between two paths from p to f , that is, no one of the two paths preempts the other, and so none of them can be considered as valid, both in Horty’s and in our approach. In Horty’s interpretation, such a conflict prevents also the construction of paths starting at p and passing through f : in order to be constructible a path has to be built augmenting a valid shorter path, and thus we cannot construct any path starting from p and passing through f (Horty, 1994, Def. 2.1.1). So, an unresolved conflict totally eliminates the possibility to consider such paths in more ample argumentations, where they could play some role. On the other hand, in our approach we are not so radical about conflicts: the fact that we cannot conclude neither $p \rightarrow f$ nor $p \not\rightarrow f$ does not eliminate the possibility that in the actual world one of such connections is true; simply we do not have enough information to decide. The possibility of $p \rightarrow f$ to be effectively valid invites us to take under consideration such a potential argumentation in moving from p to b . So, looking for a connection between p and b in Example 3.4, while Horty cannot consider the path $\langle p, c, t, f, b \rangle$, avoiding the rise of a conflict with the path $\langle p, e, b \rangle$, in our approach we still consider the possibility for $\langle p, c, t, f \rangle$ to be effectively true, allowing the path $\langle p, c, t, f, b \rangle$ to play a role in deciding whether there is a valid connection between p and b . In such a way we have a potential conflict with $\langle p, e, b \rangle$ that prevents the validity of the latter. For other significant examples of our approach see Appendix A.

Notice that, even if it is built on the notion of courses, our procedure ‘respects’ the classical notion of potential path, that is, every valid connection corresponds to a potential path on the net (Definition 2.1).

Proposition 3.2. *Consider a net \mathcal{N} . For every connection $p \vdash_{\mathcal{N}} q$ (resp., $p \vdash_{\mathcal{N}} \neg q$) validated by our procedure, there is a corresponding positive (resp., negative) potential path from p to q in the net \mathcal{N} .*

3.1.3 INFERENCE RELATION

Talking about nets, the structural properties characterizing rational consequence relations, *REF*, *CT*, *CM*, and *RM*, take the following form⁵:

$$\begin{array}{l}
 \text{(REF)} \quad \mathcal{N} \rightsquigarrow p \star q \text{ for every } p \star q \in \mathcal{N} \\
 \text{(CT)} \quad \frac{\mathcal{N} \rightsquigarrow p \star q \quad \mathcal{N}, p \star q \rightsquigarrow r \star s}{\mathcal{N} \rightsquigarrow r \star s} \\
 \text{(CM)} \quad \frac{\mathcal{N} \rightsquigarrow p \star q \quad \mathcal{N} \rightsquigarrow r \star s}{\mathcal{N}, p \star q \rightsquigarrow r \star s} \\
 \text{(RM)} \quad \frac{\mathcal{N} \rightsquigarrow r \star s \quad \mathcal{N} \not\rightarrow p \star q}{\mathcal{N}, p \not\star q \rightsquigarrow r \star s}
 \end{array}$$

The meaning of the properties is still the same as the propositional case, simply readapted to the expressivity of the INs: the net represents the information at our disposal, the premises of the derivation, and the links are the informational atoms of our language.

5. with $\star, \ast \in \{\rightarrow, \not\rightarrow\}$; ‘ $\mathcal{N}, a \star b$ ’ in the premises indicates the addition of the direct link $a \star b$ to the net \mathcal{N} ; $\not\star$ indicates the opposite arrow of \star (e.g. $\not\rightarrow$ is \rightarrow iff \star is $\not\rightarrow$)

Hence, the sense of the rules is the same as before. (REF) indicates that whatever piece of information (link) is in the premises, it appears also in the conclusions. (CT) is a cut condition, that states that if the validity of a link can be derived from the links in the rest of the net, such a link can be eliminated without affecting the set of the conclusions derivable from the net. (CM) is a form of constrained monotony, opposite to (CT), that states that whatever conclusion can be derived from the net, it can be added to the premises without affecting the other conclusions. (CT) and (CM) have an intuitive appeal, and from the logical point of view characterize \rightsquigarrow as a closure operation. The translation of (RM) is less intuitive, since we do not have in INs a classical notion of negation, but we have only a notion of conflict; hence the sense of the rule is that, if $p \rightarrow q$ is not a consequence of \mathcal{N} , then the addition of information conflicting with $p \rightarrow q$, *i.e.*, $p \nrightarrow q$, should not affect the defeasible consequences of the net \mathcal{N} . The fact that INs do not share with classical logic the notions of contradiction and negation makes this formulation of (RM) less intuitive and interesting.

The proprieties (REF), (CT) and (CM) are often considered proprieties that a non-monotonic consequence relation should satisfy (Kraus et al., 1990; Makinson, 1994), and so it is interesting to check if they are satisfied in the IN formalism. We know that the classical approaches to inheritance nets do not satisfy (CT) and (CM) (Makinson, 1994, pp. 56-57), while our approach is logically more appealing.

Proposition 3.3. \rightsquigarrow' satisfies (REF), (CT) and (CM).

While (RM) is not satisfied.

Proposition 3.4. \rightsquigarrow' does not satisfy (RM).

The following example proves the proposition

Example 3.5. Consider the net in Figure 5. The net is composed of the links $p \rightarrow f$, $f \rightarrow b$, and $p \rightarrow t$. We have $\mathcal{N} \rightsquigarrow p \rightarrow b$ and $\mathcal{N} \nrightarrow t \rightarrow b$, but $\mathcal{N}, t \nrightarrow b \nrightarrow p \rightarrow b$.




Figure 5: Counterexample to RM.

This example actually shows that, dealing with the notion of negation and consistency that characterize INs, (RM) does not look as a desirable property anymore, since the addition of $t \nrightarrow b$ to the net creates a *Nixon Diamond* from which we do not want to derive $p \rightarrow b$ (see Example 3.2).

Other properties of logical consequence relations, *left equivalence* and *right weakening*, have an analogous in the following properties:

$$(LE) \quad \frac{\mathcal{N} \rightsquigarrow p \star q \quad \mathcal{N} \rightsquigarrow p \Rightarrow r \quad \mathcal{N} \rightsquigarrow r \Rightarrow p}{\mathcal{N} \rightsquigarrow r \star q}$$

$$(RW) \quad \frac{\mathcal{N} \rightsquigarrow p \rightarrow q \quad \mathcal{N} \rightsquigarrow q \Rightarrow r}{\mathcal{N} \rightsquigarrow p \rightarrow r}$$

We also introduce a property that corresponds to the logical property of *supra classicality* (if $C \models D$, then $C \vdash D$), a rule satisfied by most non-monotonic consequence relations:

$$(Sup) \quad \frac{\mathcal{N} \rightsquigarrow p \Rightarrow q}{\mathcal{N} \rightsquigarrow p \rightarrow q} \quad \frac{\mathcal{N} \rightsquigarrow p \not\Leftarrow q}{\mathcal{N} \rightsquigarrow p \not\rightarrow q}$$

Proposition 3.5. \rightsquigarrow' satisfies (LE), (RW), and (Sup).

3.1.4 CONSISTENCY

As indicated at the end of Section 2.1, a net is considered inconsistent if we are forced to conclude for some pair of nodes p, q that p and q are both positively and negatively connected. Since, as seen above, \rightsquigarrow' satisfies (Sup), we can say that a mixed net is consistent iff we cannot conclude both $\mathcal{N} \rightsquigarrow' p \rightarrow q$ and $\mathcal{N} \rightsquigarrow' p \not\rightarrow q$ for any pair of nodes p, q .

Now we are going to see that in order to check the consistency of a mixed net we can use the ranking procedure: it is sufficient to apply it to the whole net. As for the propositional case (see Section 2.2), the ranking procedure defined on the nodes of a net terminates, after a finite number of steps, into either an empty set α_n or a fixed point of the function E , i.e., the set of the nodes that result always exceptional. In such a case, we say that such nodes have an infinite ranking value ($\tau(p) = \infty$). If we want to check whether a net \mathcal{N} is consistent, it is sufficient to apply the ranking procedure to the entire net, and see if there are nodes with infinite ranking.

Proposition 3.6. A net \mathcal{N} is consistent iff we do not have a node p with $\tau(p) = \infty$, that is, we do not conclude both $\mathcal{N} \rightsquigarrow' p \rightarrow q$ and $\mathcal{N} \rightsquigarrow' p \not\rightarrow q$ for any pair p, q .

Example 3.6. The net in Figure 6 is an example of an inconsistent net, from which we would conclude both $\mathcal{N} \rightsquigarrow' t \rightarrow f$ and $\mathcal{N} \rightsquigarrow' t \not\rightarrow f$. Such a net is translated into the




Figure 6: An example of inconsistent net.

knowledge base

$$\begin{aligned}\Phi &= \{t \supset b, t \supset p, p \supset b, b \supset p\} \\ \Delta &= \{p \supset \neg f, b \supset f\}.\end{aligned}$$

We proceed the ranking of the entire net, and we obtain that $\Delta \models_{\Phi} \neg p$, $\Delta \models_{\Phi} \neg b$ and $\Delta \models_{\Phi} \neg t$, that is, $\mathcal{E}_1 = \Delta$. Hence, Δ is a fixed-point of the exceptionality ranking function, and p, b, t have ∞ as ranking value.

3.1.5 PROPERTIES

In the field of inheritance networks, a taxonomy of the different approaches has been developed on the basis of some relevant properties (Horty, 1994). We briefly check which of them are satisfied by our approach.

- *Purely defeasible / mixed nets. Cyclic / acyclic nets.* Our procedure deals easily with two properties that often create problems in the traditional approaches: the presence of both strict and defeasible links (*mixed nets*), and the presence of cycles (*cyclic nets*).
- *Credulous / skeptical / directly skeptical approaches.* Our approach corresponds to a *directly skeptical* approach: given a net, we obtain a unique set of valid connections in it (vs. the credulous approach, that allows for different sets of valid paths, possibly in conflict with each other), and such a unique set is not obtained from the intersection of different possible extensions (as in some skeptical approaches), but it is obtained from a single closure operation.
- *Upward / downward chaining.* For the definition of valid paths, we do not use any form of induction on their length, neither starting from the initial node toward the terminal node (*upward chaining*), nor in the reverse direction (*downward chaining*); hence, no form of chaining is used in our procedure.
- *On-path / off-path preemption.* Off-path preemption is the classical form of preemption, used also by Horty (1994, Def. 3.2.2), while on-path preemption is more binding, requiring the preempting node to lie on the initial segment of the path it preempts (Horty, 1994, sect. 4.2.4). We do not exactly formalise a form of preemption, since we do not confront directly the different paths between two nodes. However, our procedure has a behaviour that is analogous to the use of off-path preemption.

3.1.6 COMPUTATIONAL COMPLEXITY

To define the overall complexity of our decision procedure over the nets, we have to consider the complexity of the course-identification procedure, that is, given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$ and two nodes s, t in \mathcal{N} , which is the computational cost to identify $\Delta_{s,t}$ (note that Φ can easily be computed in polynomial time), whose size is bounded polynomially by the size of \mathcal{N} . Given that the construction of courses is independent from the nature of the links (either they are positive or negative, defeasible or strict), we can analyse the problem using simple directed graphs. Given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, it is sufficient to define the correspondent directed graph $G = \langle V, E \rangle$ in the following way:

- V is the set of nodes in \mathcal{N} .
- E is a set of directed links $\langle a, b \rangle$, with $a, b \in V$, s.t. $\langle a, b \rangle \in E$ iff one of the following holds:

$$a \Rightarrow b \in \mathcal{S} , \quad a \not\Leftarrow b \in \mathcal{S} , \quad a \rightarrow b \in \mathcal{U} , \quad a \not\Leftarrow b \in \mathcal{U} .$$

Recall that we have stated the presence of $a \not\Leftarrow b$ in \mathcal{S} implies that $b \not\Leftarrow a$ is in \mathcal{S} too. So, for $a \not\Leftarrow b \in \mathcal{S}$ we have both $\langle a, b \rangle$ and $\langle b, a \rangle$ in E .

Once we have defined G , let us recall a well-known result in graph theory saying that in a directed graph, given two nodes p and q , determining if there is a path from p to q can be determined in time $\mathcal{O}(|V| + |E|)$, *e.g.* using BFS (Breadth First Search) (Cormen, Stein, Rivest, & Leiserson, 2001). Now, the following argument shows that indeed $\Delta_{s,t}$ can be determined in polynomial time⁶. At first, we check if there is a path between s and t . If not, then $\Delta_{s,t} = \emptyset$. Otherwise, we call the procedure **Delta**(s) below:

Delta(s): for each outgoing edge $\langle s, x \rangle$ of s , such that both $\langle s, x \rangle$ and x are not marked, do: if there is a path between x and t then mark both $\langle s, x \rangle$ and x , and recursively, call **Delta**(x).

Once finished, $\Delta_{s,t}$ can immediately be build from the marked edges. Note that each edge is marked once and each node is marked (*i.e.*, explored) once and, thus, the algorithm is bounded polynomially by the size of the graph.

Once we have found the set $\Delta_{s,t}$ and Φ , we have to apply the decision procedure based on the propositional rational closure to decide if there is a valid connection between p and q (as for Section 2.2, the number of entailment tests is polynomially bounded by the size of the net). As all formulae are 2-CNF, like Proposition 3.1, we obtain that the decision procedure w.r.t. the net respects the complexity costs of the related propositional calculus.

Proposition 3.7. *Deciding defeasible consequence under inheritance networks (\rightsquigarrow') can be done in polynomial time.*

Eventually, if we want to determine all the valid links in the net \mathcal{N} we have to consider all the pairs of nodes in the net \mathcal{N} . So, obtained the graph $G = \langle V, E \rangle$ we have to repeat the procedure for all the elements of the set of the pairs of nodes in the graph, whose cardinality is $|V|(|V| - 1)$. Hence, again

Proposition 3.8. *Computing all the valid connections in a net can be done in polynomial time.*

3.2 Boolean Inheritance Nets

We next extend INs by introducing in them the classical propositional connectives \neg, \wedge, \vee . Despite such an extension has been felt as desirable, we are aware of just an attempt in this direction (Horty & Thomason, 1990).

6. We are not interested here in figuring out a tight bound.




Figure 7: Disjunction




Figure 8: Conjunction

3.2.1 NEGATION

So far, we have used the link $\not\Leftarrow$ to indicate that two classes are disjoint: $p \not\Leftarrow q$ has $p \supset \neg q$ as logical meaning. We change the notation and substitute $\not\Leftarrow$ with \Leftrightarrow^\neg , indicating with $p \Leftrightarrow^\neg q$ that ‘class p and class q are complementary’ (i.e., $p \equiv \neg q$), and in general we will indicate the complementary class of a class p with \bar{p} . Hence, we can substitute every link $p \not\Leftarrow q$ in a net with four links: $p \Leftrightarrow^\neg \bar{p}$, $q \Leftrightarrow^\neg \bar{q}$, $p \Rightarrow \bar{q}$, and $q \Rightarrow \bar{p}$. Moreover, we can eliminate the negative defeasible links, since $p \not\Leftarrow q$ can be expressed as $p \rightarrow \bar{q} \Leftrightarrow^\neg q$. So, we can transform an IN into a net using only the arrows \rightarrow , \Rightarrow , and \Leftrightarrow^\neg . We shall continue to use $\not\Leftarrow$ as a macro indicating the valid negative strict connections obtained from the composition of \Leftrightarrow^\neg and \Rightarrow , that is, we indicate with $p \not\Leftarrow q$ the presence of a path

$$p \underbrace{\Rightarrow a \Rightarrow \dots \Rightarrow b}_{n \text{ arrows}} \Leftrightarrow^\neg \underbrace{c \Leftarrow \dots \Leftarrow d}_{m \text{ arrows}} q,$$

with $n, m \geq 0$.

3.2.2 CONJUNCTION AND DISJUNCTION

Next, we extend inheritance nets to support conjunction and disjunction as well, by allowing links $a, b \Leftrightarrow^\wedge c$ (conjunction of a and b is equivalent to c) and $c \Leftrightarrow^\vee a, b$ (disjunction of a and b is equivalent to c). We will assume that inheritance nets containing such kind of links are closed according to the following rule: if there is $a, b \Leftrightarrow^\wedge c$ (resp., $a, b \Leftrightarrow^\vee c$) in a net, then there are also $c \Rightarrow a$ and $c \Rightarrow b$ (resp., $a \Rightarrow c$ and $b \Rightarrow c$) in the net. We call these nets *Boolean Defeasible Inheritance Networks* (BINs). We shall use $a \otimes b$ and $a \oplus b$ to indicate, respectively, that a node represents the conjunction or the disjunction of a and b . Graphically, we indicate disjunctive and conjunctive links as in Figure 7 and Figure 8, respectively.

We extend now our reasoning method to BINs. To do so, we need to amplify the notion of *course*, introducing the notion of *duct*: we consider not only ‘linear’ routes from one point to another, but also ‘parallel’ routes, in order to model the introduction of the conjunction in the consequent and the introduction of the disjunction in the antecedent. Roughly,

$$\pi = \langle s, \frac{\sigma}{\sigma'}, t \rangle$$

will indicate a duct π that starts at node s and develops through the ducts σ and σ' , both reaching the node t .

Definition 3.2 (Duct). Ducts are defined as follows (where $\star \in \{\Rightarrow, \Leftrightarrow, \rightarrow, \nrightarrow\}$):

1. every link $p \star q$ in \mathcal{N} is a duct $\pi = \langle p, q \rangle$ in \mathcal{N} ;
2. if $\pi = \langle \sigma, q \rangle$ is a duct and $q \star r$ is a link in \mathcal{N} that does not already appear in π , then $\pi' = \langle \pi, r \rangle$ is a duct in \mathcal{N} ;
3. if $\pi = \langle q, \sigma \rangle$ is a duct and $r \star q$ is a link in \mathcal{N} that does not already appear in π , then $\pi' = \langle r, \pi \rangle$ is a duct in \mathcal{N} ;
4. if $\langle t, \sigma, p \rangle$ and $\langle r, \sigma', p \rangle$ are ducts, then for $s \Leftrightarrow^\vee t, r \in \mathcal{S}$, $\langle s, \frac{t, \sigma}{r, \sigma'}, p \rangle$ is a duct; and
5. if $\langle p, \sigma, t \rangle$ and $\langle p, \sigma', r \rangle$ are ducts, then for $t, r \Leftrightarrow^\wedge s \in \mathcal{S}$, $\langle p, \frac{\sigma, t}{\sigma', r}, s \rangle$ is a duct.

Now our reasoning method for BINs is as follows. Given a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$, we can define a correspondent knowledge base $\mathcal{K} = \langle \Phi, \Delta \rangle$, where

$$\begin{aligned} \Phi &= \{p \supset q \mid p \Rightarrow q \in \mathcal{S}\} \\ &\cup \{p \equiv \neg q \mid p \Leftrightarrow^\neg q \in \mathcal{S}\} \\ &\cup \{p \equiv q \wedge r \mid q, r \Leftrightarrow^\wedge p \in \mathcal{S}\} \\ &\cup \{p \equiv q \vee r \mid p \Leftrightarrow^\vee q, r \in \mathcal{S}\} \\ &\cup \{p \supset \neg q \mid p \nrightarrow q \in \mathcal{S}\} \end{aligned}$$

$$\Delta = \{p \supset q \mid p \rightarrow q \in \mathcal{U}\} .$$

Now, we may proceed to the definition of $\sim_{\mathcal{N}}$ as in Section 3.1, simply considering $\mathcal{C}^{\mathcal{N}}$ as the set of the ducts in \mathcal{N} , and $\mathcal{C}_{p,q}^{\mathcal{N}}$ (or simply $\mathcal{C}_{p,q}$) as the set of the ducts from p to q .

Example 3.7. Consider the net \mathcal{N} illustrated in Figure 9. The net \mathcal{N} is mapped into the




Figure 9: Example 3.7

KB $\mathcal{K} = \langle \Phi, \Delta \rangle$, where ⁷

$$\begin{aligned} \Phi &= \{c \wedge d \equiv g, f \equiv \neg g\} \\ \Delta &= \{a \supset b, b \supset c, b \supset d, b \supset e, a \supset f\} . \end{aligned}$$

7. To ease the reading, we have omitted the redundant implications such as $g \supset c$, obtained from $c, d \Leftrightarrow^\wedge g, g \Rightarrow c \in \mathcal{N}$.

Now, we ask whether a is connected to c . It can be verified that

$$\Delta_{a,c} = \{a \supset b, b \supset c, b \supset d, a \supset f\} .$$

Note that $b \supset d \in \Delta_{a,c}$, as there is a duct from a to c that passes through c and d in order to reach g , and then back towards c . Now, the only negated antecedent is a ($\Delta_{a,c} \models_{\Phi} \neg a$) and, thus,

$$\widehat{\Delta}_{a,c} = \{a \supset b, a \supset f\} .$$

Since $\widehat{\Delta}_{a,c} \not\models_{\Phi} a \supset c$ and $\widehat{\Delta}_{a,c} \not\models_{\Phi} a \supset \neg c$, we have

$$a \not\sim_{\mathcal{N}} c \text{ and } a \not\sim_{\mathcal{N}} \neg c .$$

In a similar way, we may show that $a \not\sim_{\mathcal{N}} d$ and $a \not\sim_{\mathcal{N}} \neg d$. This is the desirable result: since $a \rightarrow f$ is a direct link, we have that $a \sim_{\mathcal{N}} f$ (i.e., $a \sim_{\mathcal{N}} \neg(c \wedge d)$), and hence we know that we cannot conclude both $a \sim_{\mathcal{N}} c$ and $a \sim_{\mathcal{N}} d$. But, since we have no evidence whether one of such conclusions has to be preferred to the other, we do not conclude either of them. The result of our ‘skeptical’ approach is that $a \not\sim_{\mathcal{N}} c$, $a \not\sim_{\mathcal{N}} \neg c$, $a \not\sim_{\mathcal{N}} d$, and $a \not\sim_{\mathcal{N}} \neg d$. On the other hand, since the only duct connecting a to e is $\langle a, b, e \rangle$ (that is, the nodes c , d , g , and f do not play any role in any possible argumentation connecting a to e), we can conclude $a \sim_{\mathcal{N}} e$.

3.2.3 PROPERTIES

We call $\rightsquigarrow_{\text{BIN}}$ the inference relation defined by the just defined closure operation over BINs and we have:

$$\begin{aligned} \mathcal{N} \rightsquigarrow_{\text{BIN}} p \Rightarrow q & \text{ iff } p \models_{\Phi} q \\ \mathcal{N} \rightsquigarrow_{\text{BIN}} p \not\Leftarrow q & \text{ iff } p \models_{\Phi} \neg q \\ \mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow q & \text{ iff } p \sim_{\mathcal{N}} q \\ \mathcal{N} \rightsquigarrow_{\text{BIN}} p \not\rightarrow q & \text{ iff } p \sim_{\mathcal{N}} \neg q . \end{aligned}$$

BINs inherit the same structural properties of our INs, that is, (REF), (CM), and (CT). Analogously, (LE), (RW), and (Sup) are still valid.

Proposition 3.9. $\rightsquigarrow_{\text{BIN}}$ satisfies (REF), (CM), and (CT).

Proposition 3.10. $\rightsquigarrow_{\text{BIN}}$ satisfies (LE), (RW), and (Sup).

Since the procedure defined in Section 3.1 is simply a special case of the procedure for BINs, (RM) is falsified also for BINs by the same counter-example of proposition 3.4. Moreover, introduced conjunction and disjunction, we can express the analogous of the rules of disjunction in the premises (OR) and conjunction in the consequent (AND):

$$\begin{aligned} \text{(OR)} \quad & \frac{\mathcal{N} \rightsquigarrow_{\text{BIN}} p \star q \quad \mathcal{N} \rightsquigarrow_{\text{BIN}} s \star q \quad \mathcal{N} \rightsquigarrow_{\text{BIN}} t \Leftrightarrow^{\vee} p, s}{\mathcal{N} \rightsquigarrow_{\text{BIN}} t \star q} \\ \text{(AND)} \quad & \frac{\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow q \quad \mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow s \quad \mathcal{N} \rightsquigarrow_{\text{BIN}} q, s \Leftrightarrow^{\wedge} t}{\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow t} \end{aligned}$$

The sense is the same as in the propositional case, and it remains intuitive also in the BIN environment: (OR) represents the validity of reasoning by cases, while (AND) represents that the conjunction of distinct conclusions is still a valid conclusion from the net.

Proposition 3.11. $\rightsquigarrow_{\text{BIN}}$ satisfies (OR) and (AND).

3.2.4 CONSISTENCY

Also w.r.t. consistency, we obtain the same result as for INs, *i.e.*, the net is consistent if the ranking procedure terminates into an empty set.

Proposition 3.12. A BIN \mathcal{N} is consistent iff we do not have any node p with $\mathfrak{r}(p) = \infty$, that is, we cannot conclude both $p \rightarrow q$ and $p \not\rightarrow q$ for any pair p, q .

Remark 2. As seen in Section 3.1.4 the notion of consistency in inheritance nets is different from the notion of consistency for propositional logic. Using our procedure a net is inconsistent if, applying the ranking function to the entire net, we obtain a node with ∞ as ranking value (see Proposition 3.6). Up to this point, if we find that a net is inconsistent, we simply stop our decision procedure.

In what follows we are going to work with BINs in the framework of propositional logic. So, in order to assimilate the notion of consistency with the one of propositional logic, from now on we shall consider a modified version of our procedure for BINs. Suppose we have to decide the validity of the connection between two nodes p, q in a net \mathcal{N} . If \mathcal{N} results consistent, then we proceed as above, otherwise, if our net results inconsistent (some node has infinite ranking value) we do not simply stop, but, in case $\mathfrak{r}(p) < \infty$, we still apply our procedure. Otherwise, if p itself has infinite ranking value ($\mathfrak{r}(p) = \infty$), we do not proceed further.

3.2.5 COMPUTATIONAL COMPLEXITY

As for INs, we have to determine $\Delta_{s,t}$ (computing Φ is again immediate) from the ducts in a BIN. Now, it is not difficult to see that a recursive BFS graph travelling procedure as the one devised for INs can be worked out for BINs as well. By illustration, refer to Figure 9 and assume we are processing node b . Since from b , both d and c are reachable and $d, c \Leftrightarrow^{\vee} g \in \mathcal{S}$, and, recursively, there is a path from g to c , we can mark $\langle b, c \rangle$ and $\langle b, d \rangle$, and mark the “conjunction” $d, c \Leftrightarrow^{\vee} g \in \mathcal{S}$ as visited. Again, all nodes and “aggregated nodes” are visited at most ones, guaranteeing polynomial cost for computing $\Delta_{s,t}$.

Now, while $\Delta_{s,t}$ and Φ can be determined in polynomial time, and as for Section 2.2, the number of entailment tests is polynomially bounded by the size of the net, the strict part may encode any propositional formula and, thus, unlike the case of INs, we have:

Proposition 3.13. Deciding defeasible consequence for BINs ($\rightsquigarrow_{\text{BIN}}$) is a coNP-complete problem.

4. Defeasible Inheritance in Propositional Logic

Now, we depart from BINs and apply a similar reasoning procedure in the framework of propositional logic, and show how to obtain a kind of closure of a knowledge base that

results to be a rational consequence relation, but that is more informative than the classical rational closure (Lehmann & Magidor, 1992).

We consider a propositional language with only \neg, \vee, \wedge as connectives. So, we start with a conditional KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ (see Section 2.2), with $\mathcal{T} = \{C_1 \vdash D_1, \dots, C_n \vdash D_n\}$ and $\mathcal{D} = \{E_1 \sim F_1, \dots, E_m \sim F_m\}$.

Step 1. Given a conditional base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, check if \mathcal{K} is preferentially consistent (that is, check if its materialisation is consistent; Section 2.2, Step 4.1). If it is consistent, we define a BIN from \mathcal{K} , *i.e.*, a net $\mathcal{N}_{\mathcal{K}} = \langle \mathcal{S}_{\mathcal{K}}, \mathcal{U}_{\mathcal{K}} \rangle$, modelling the information in \mathcal{K} in the following way:

- (i) we consider every formula C that appears as antecedent or as succedent in the conditionals in \mathcal{K} , and we create a node C representing each of them, modulo logical equivalence (that is, a node C represents the class of the formulae logically equivalent to C).
- (ii) for each node we add also, if not already present, the complementary node (the node representing its negation), and we link them by $\Leftrightarrow \neg$;
- (iii) we add the strict links: if $C \vdash D \in \mathcal{T}$ we add the strict link $C \Rightarrow D$ to the net; we also add to $\mathcal{S}_{\mathcal{K}}$ all the strict links that correspond to the logical dependencies between the formulae represented by the nodes w.r.t. the consequence relation $\models_{\mathcal{T}}$. If \top appears in a conditional, we add to the net a correspondent node \top , and, for every other node n in the net, we add a strict node $n \Rightarrow \top$. Analogously, for \perp we add $\perp \Rightarrow n$ for every node n in the net.
- (iv) eventually, for $C \sim D \in \mathcal{D}$, we add a defeasible link \rightarrow from node C to node D .

Step 2. We apply the reasoning procedure for BINs to $\mathcal{N}_{\mathcal{K}}$ (Section 3.2) to identify all valid defeasible connections $C \sim_{\mathcal{N}} D$ and we add them as $C \sim D$ to the conditional base \mathcal{K} to obtain a new conditional base $\mathcal{K}' = \langle \mathcal{T}, \mathcal{D}' \rangle$ ⁸.

Step 3. Finally, we apply to \mathcal{K}' its rational closure (Section 2.2) and we define a non-monotonic consequence relation $\vdash_{\mathcal{K}}$ by

$$C \vdash_{\mathcal{K}} D \text{ iff } C \sim D \in \mathbb{R}(\mathcal{K}').$$

Now, we can show that

Proposition 4.1. $\vdash_{\mathcal{K}}$ is a rational consequence relation containing \mathcal{K} .

Example 4.1. Consider again the penguin example. We modify it slightly in order to consider also the use of the connectives. While birds (b) typically fly, live on trees, and have wings (f, t, w), penguins do not fly and do not live on trees ($\neg f \wedge \neg t$). So, our knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ will be:

$$\begin{aligned} \mathcal{T} &= \{p \vdash b\} \\ \mathcal{D} &= \{b \sim f, b \sim t, b \sim w, p \sim \neg f \wedge \neg t\}. \end{aligned}$$

8. We do not modify \mathcal{T} , since all the strict connections valid in the net are classically derivable from \mathcal{T} .

Notwithstanding penguins' atypicality as birds, penguins have wings, and we would like to be able to derive it from the information at our disposal, that is, we would like to conclude $p \vdash w$. Please note that this is not possible using classical preferential approaches, but we can obtain such a conclusion passing through the first step of our closure operation, that is, defining the corresponding net.

Specifically, from the knowledge base \mathcal{K} we define the net in Figure 10 (the dashed arrows are the strict arrows that are not explicit in the conditional base, but that are logically valid and that are added to $\mathcal{S}_{\mathcal{K}}$ in the construction of the net $\mathcal{N}_{\mathcal{K}}$).




Figure 10: Example 4.1

Using the procedure defined for BINs, from such a net we obtain a new knowledge base

$$\mathcal{K}' = \langle \mathcal{T}, \mathcal{D}' \rangle ,$$

with

$$\mathcal{D}' = \{ b \vdash f, b \vdash t, b \vdash w, p \vdash \neg f \wedge \neg t, p \vdash w, p \vdash \neg f, p \vdash \neg t, b \vdash f \vee t \} .$$

Note that, while the new conditionals

$$p \vdash \neg f, \quad p \vdash \neg t, \quad b \vdash f \vee t$$

would be present also in the simple rational closure of \mathcal{K} (we obtain them by Right Weakening), we have obtained also the conditional

$$p \vdash w ,$$

that would not be present in the rational closure of \mathcal{K} (see Remark 1).

Now, following the procedure defined in Section 2.2, we compute the rational closure of the new knowledge base \mathcal{K}' , obtaining a rational consequence relation that contains the original \mathcal{K} .

Please note that, if we were using only BINs, we could have not derived anything else, since our vocabulary would be limited to the propositions expressed by the nodes; however, by relying on the rational closure of propositional knowledge bases, we can reason using the full expressivity of a propositional language, deriving new conditionals as, for example,

$$b \wedge \text{green} \sim f ,$$

which can not be derived using BINs as *green* does not appear in the net.

The next example shows another characteristic of this approach. In the preferential approach typicality is an ‘absolute’ property of a proposition w.r.t. the agent’s knowledge base, that is, if a class results atypical w.r.t. some other class (as penguins w.r.t. birds), it results atypical w.r.t. the entire knowledge base. In our approach instead, typicality is a comparative notion: we can consider a class as exceptional with respect to a superclass, but absolutely typical with respect to another.

Example 4.2. Consider a red fish (r). It is both a fish (f) and a pet (p). Typically, a fish has gills (g) and scales (s), while pets are docile (d) and play with kids (k). Red fishes are not typical pets, since they do not play with kids. So, the $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ is

$$\begin{aligned} \mathcal{T} &= \{r \vdash f, r \vdash p\} \\ \mathcal{D} &= \{r \sim \neg k, p \vdash k, p \vdash d, f \vdash g, f \vdash s\} . \end{aligned}$$

In the rational closure red fishes, since they are atypical pets (they do not play with kids), result atypical in general, and they cannot inherit any of the typical properties of all their super classes.

Instead, we want red fishes to inherit, besides the properties of pets that are compatible with them (d), also all the typical properties of fishes (g and s), since we consider them typical fishes.

To do so, we translate our knowledge base into the net in Figure 11. From such a net we obtain a new knowledge base $\mathcal{K}' = \langle \mathcal{T}, \mathcal{D}' \rangle$, with

$$\mathcal{D}' = \{r \sim \neg k, p \vdash k, p \vdash d, f \vdash g, f \vdash s, r \vdash d, r \vdash g, r \vdash s\}$$

and so we have derived exactly the desired conditionals.

Next, we compute the rational closure of \mathcal{K}' , following the procedure defined in Section 2.2, and we obtain a rational consequence relation containing \mathcal{K} and more information about red fishes, information that, as intuitive as it is, we would not be able to derive from the simple rational closure of \mathcal{K} .

Therefore, we have defined a new rational consequence relation for \mathcal{K} that extends \mathcal{K} , as $\mathcal{K} \subset \mathbb{R}(\mathcal{K}')$, and that contains intuitive conditionals not in the rational closure of \mathcal{K} .




Figure 11: Example 4.2

Consistency. Defined some inference procedure, a conditional base \mathcal{K} is consistent iff we cannot derive $\top \vdash \perp$. We have seen (Section 2.2) that for rational closure a conditional base is consistent iff its preferential closure is consistent ($\top \vdash \perp \in \mathbb{R}(\mathcal{K})$ iff $\top \vdash \perp \in \mathbb{P}(\mathcal{K})$). Here, given a base \mathcal{K} , we obtain that our procedure preserves the preferential consistency of \mathcal{K} : as seen in Section 2.2, \mathcal{K} is preferentially consistent iff its rational closure is consistent (*i.e.*, $\top \not\vdash_{\mathcal{K}^{rc}} \perp$), and we can prove the following.

Proposition 4.2. *Given a conditional base \mathcal{K} , $\top \vdash_{\mathcal{K}} \perp$ iff $\top \vdash \perp \in \mathbb{R}(\mathcal{K})$.*

By the results in Section 2.2, this corresponds to saying that $\top \vdash_{\mathcal{K}} \perp$ iff $\overline{\mathcal{K}} \models \perp$ (assuming $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, $\overline{\mathcal{K}} = \overline{\mathcal{T}} \cup \overline{\mathcal{D}}$).

Computational Complexity. Considering the procedures defined for the BINs, we can conclude that the defined procedure has the same complexity of the rational closure, as it is the composition of the procedure defined for BINs (Proposition 3.13) and a final rational closure operation (Proposition 2.2).

Proposition 4.3. *For a propositional conditional base \mathcal{K} , deciding $C \vdash_{\mathcal{K}} D$ is a coNP-complete problem.*

5. Defeasible Inheritance in DLs

Next, we apply our method to a significant DL representative, namely \mathcal{ALC} (Baader et al., 2003, ch. 2). \mathcal{ALC} corresponds to a fragment of first order logic, using monadic predicates, called *concepts*, and dyadic ones, called *roles*.

In order to stress the parallel between the procedure presented in Section 2.2 and the proposal in \mathcal{ALC} , we are going to use the same notation for the components playing an analogous role in the two construction: we use p, q, r, \dots for *concept names*, and C, D, E, \dots to indicate *concepts* in general, instead, respectively, of atomic propositions and propositions, and \models and \vdash to indicate, respectively, the classical consequence relation of \mathcal{ALC} and

a non-monotonic consequence relation in \mathcal{ALC} . δ will indicate a *default concept*, that is, a concept that we assume as applying to every individual, if not informed of the contrary.

We have a finite set of *concept names* $\mathcal{C} = \{p, q, r, \dots\}$, a finite set of *role names* $\mathcal{R} = \{R, S, T, \dots\}$ and the set \mathcal{L} of \mathcal{ALC} -concepts is defined inductively as follows: (i) $\mathcal{C} \subset \mathcal{L}$; (ii) $\top, \perp \in \mathcal{L}$; (iii) $C, D \in \mathcal{L} \Rightarrow C \sqcap D, C \sqcup D, \neg C \in \mathcal{L}$; and (iv) $C \in \mathcal{L}, R \in \mathcal{R} \Rightarrow \exists R.C, \forall R.C \in \mathcal{L}$. Concept $C \supset D$ is used as a shortcut of $\neg C \sqcup D$. The symbols \sqcap and \sqcup correspond, respectively, to the conjunction \wedge and the disjunction \vee of classical logic.

Given a set of *individuals* \mathcal{O} , indicated by bold letters $\mathbf{a}, \mathbf{b}, \mathbf{c}, \dots$, an *assertion* is of the form $\mathbf{a}:C$ ($C \in \mathcal{L}$) or of the form $(\mathbf{a}, \mathbf{b}):R$ ($R \in \mathcal{R}$), respectively indicating that the individual \mathbf{a} is an instance of concept C , and that the individuals \mathbf{a} and \mathbf{b} are connected by the role R .

A *general inclusion axiom* (GCI) is of the form $C \sqsubseteq D$ ($C, D \in \mathcal{L}$) and indicates that any instance of C is also an instance of D . We use $C = D$ as a shortcut of the pair of $C \sqsubseteq D$ and $D \sqsubseteq C$.

From a FOL point of view, assertions and inclusion axioms can easily be mapped in FOL by the following transformation:

$$\begin{array}{ll} \tau(\mathbf{a}:C) = \tau(\mathbf{a}, C), & \tau((\mathbf{a}, \mathbf{b}):R) = R(\mathbf{a}, \mathbf{b}), \\ \tau(C \sqsubseteq D) = \forall x.(\tau(x, C) \rightarrow \tau(x, D)), & \tau(x, \top) = \top(x), \quad \tau(x, \perp) = \perp(x), \\ \tau(x, A) = A(x), & \tau(x, \neg C) = \neg\tau(x, C), \\ \tau(x, C \sqcap D) = \tau(x, C) \wedge \tau(x, D), & \tau(x, C \sqcup D) = \tau(x, C) \vee \tau(x, D), \\ \tau(x, \exists R.C) = \exists y.(R(x, y) \wedge \tau(y, C)) & \text{where } y \text{ is a new variable,} \\ \tau(x, \forall R.C) = \forall y.(R(x, y) \rightarrow \tau(y, C)) & \text{where } y \text{ is a new variable.} \end{array}$$

Now, a classical knowledge base is defined by a pair $\mathcal{K} = \langle \mathcal{A}, \mathcal{T} \rangle$, where \mathcal{T} is a finite set of GCIs (the *TBox*) and \mathcal{A} is a finite set of assertions (the *ABox*), whereas a *defeasible knowledge base* is represented by a triple $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, where additionally \mathcal{D} is a finite set of conditionals $C \sqsubset D$ ('an instance of a concept C is typically an instance of a concept D '), with $C, D \in \mathcal{L}$ ⁹.

Example 5.1. Consider again the penguin example. Just add a role *Prey* in the vocabulary, where a role instantiation $(\mathbf{a}, \mathbf{b}):Prey$ is read as ' \mathbf{a} preys on \mathbf{b} ', and add also two more concepts, *i* (insect) and *fi* (fish). An example of defeasible KB is

$$\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$$

with

$$\begin{array}{l} \mathcal{A} = \{\mathbf{a}:p, \mathbf{b}:b, (\mathbf{a}, \mathbf{c}):Prey, (\mathbf{b}, \mathbf{c}):Prey\} \\ \mathcal{T} = \{p \sqsubseteq b, i \sqsubseteq \neg fi\} \\ \mathcal{D} = \{p \sqsubset \neg f, b \sqsubset f, p \sqsubset \forall Prey.fi, b \sqsubset \forall Prey.i\}. \end{array}$$

The particular structure of a defeasible KB allows for the 'isolation' of the pair $\langle \mathcal{T}, \mathcal{D} \rangle$, that we could call the *conceptual system* of the agent, from the information about the individuals (formalized in \mathcal{A}).

9. Since for the monotonic part we substitute the meta-linguistic conditionals $C \vdash D$ with formulae $C \sqsubseteq D$, we substitute also for the defeasible part of the knowledge base the conditionals $C \vdash D$ with conditional formulae $C \sqsubset D$, that we could call *defeasible inclusion axioms*.

In what follows we are going to work with the information about concepts $\langle \mathcal{T}, \mathcal{D} \rangle$ first, exploiting the immediate analogy with the homonymous pair in the propositional setting, and then we will address the case involving individuals as well. We show that by using our method we overcome to the limits of classical rational closure, already presented for \mathcal{ALC} (Casini & Straccia, 2010), in a similar way as for the propositional case. Please note that the procedure presented here is based on a slightly modified version of the procedure for rational closure previously presented by Casini and Straccia (2010), *i.e.*, on the one presented by Britz et al. (2013). The latter is accompanied by a semantic characterization, based on DL interpretations with a preferential relation defined over the individuals. Such a semantic characterization of rational closure for \mathcal{ALC} characterises all the steps of our procedure (the ‘local’ applications of rational closure and the final one). However, we still lack a semantic characterization of the overall procedure, accounting also for the *modularization* of the knowledge base done using the INs.

Step 1. Given a conceptual system $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, check its preferential consistency, that is, define

$$\begin{aligned}\overline{\mathcal{T}} &= \{C \supset D \mid C \sqsubseteq D \in \mathcal{T}\} \\ \overline{\mathcal{D}} &= \{C \supset D \mid C \sqsubseteq D \in \mathcal{D}\}\end{aligned}$$

and construct a BIN $\mathcal{N}_{\mathcal{K}}$ from \mathcal{K} . The process is the same as the one in Section 4, just treat the concepts as propositions: nodes in $\mathcal{N}_{\mathcal{K}}$ represent the concepts appearing as antecedents or consequents of the inclusion axioms in \mathcal{T} and \mathcal{D} (modulo logical equivalence); for every node we add its complementary node, if not already present, and we connect them by \Leftrightarrow^{\neg} ; every GCI $C \sqsubseteq D \in \mathcal{T}$ becomes a strict link $C \Rightarrow D$; and every defeasible inclusion axiom $C \sqsubseteq D \in \mathcal{D}$ becomes a defeasible link $C \rightarrow D$. Moreover, consider the consequence relation $\models_{\mathcal{T}}$ as the monotonic consequence relation obtained adding the GCIs in \mathcal{T} to \models , and add to the net the strict links representing all the logical dependencies between nodes with respect to $\models_{\mathcal{T}}$ ¹⁰.

Step 2. Apply the reasoning procedure for BINs to $\mathcal{N}_{\mathcal{K}}$ (Section 3.2) to identify all the valid defeasible connections $C \rightarrow D$, and add them as $C \sqsubseteq D$ to the conditional base \mathcal{K} to obtain a new conditional base $\mathcal{K}' = \langle \mathcal{T}, \mathcal{D}' \rangle$.

Now, once we have augmented our knowledge base with new defeasible conditionals, we proceed as follows.

Step 3. Define $\mathcal{D}^* = \mathcal{D}' \cup \{C \sqcap \neg D \sqsubseteq \perp \mid C \sqsubseteq D \in \mathcal{T}\}$.

Step 4. Define $\overline{\mathcal{D}}^* = \{C \supset D \mid C \sqsubseteq D \in \mathcal{D}'\}$ and let $\mathfrak{A}_{\mathcal{D}^*} = \{C \mid C \sqsubseteq D \in \mathcal{D}^*\}$.

Step 5. Determine the exceptionality ranking of the conditionals in \mathcal{D}^* using the set of the antecedents $\mathfrak{A}_{\mathcal{D}^*}$ and the materialisations in $\overline{\mathcal{D}}^*$, where a concept C is *exceptional* w.r.t. a set of conditionals \mathcal{D} iff $\models \bigcap \overline{\mathcal{D}} \sqsubseteq \neg C$. The steps are the same of the propositional case (**Step 3** in Section 2.2) by replacing the expression $\overline{\mathcal{D}} \models \neg C$ with the expression $\models \bigcap \overline{\mathcal{D}} \sqsubseteq \neg C$. In this way define a ranking function \mathfrak{r} .

10. In order to create the strict part of the net it is possible to use the techniques introduced for the procedure of *classification* of DL knowledge bases (Baader et al., 2003, ch. 9).

Step 6. As **Step 4.1**, Section 2, verify if the KB is consistent, by checking the consistency of $\overline{\mathcal{D}^*}$. Then (as in **Steps 4.2 - 4.3** in Section 2.2), define the sets

$$\begin{aligned}\tilde{\mathcal{T}} &= \{\top \sqsubseteq \neg C \mid C \sqsubseteq D \in \mathcal{D}' \text{ and } \mathfrak{r}(C \sqsubseteq D) = \infty\} \\ \tilde{\mathcal{D}} &= \{C \sqsubseteq D \mid C \sqsubseteq D \in \mathcal{D}' \text{ and } \mathfrak{r}(C \sqsubseteq D) < \infty\}.\end{aligned}$$

Step 7. Define (similarly to **Step 5** in Section 2.2) the sets of concepts $\tilde{\Phi}$ and $\tilde{\Delta}$, with

$$\begin{aligned}\tilde{\Phi} &= \{C \mid \top \sqsubseteq C \in \mathcal{T}\} \\ \tilde{\Delta} &= \{\delta_0, \dots, \delta_n\},\end{aligned}$$

where

$$\delta_i = \prod \{C \supset D \mid C \sqsubseteq D \in \tilde{\mathcal{D}} \text{ and } \mathfrak{r}(C \sqsubseteq D) \geq i\}.$$

As for Section 2, for every δ_i , $0 \leq i < n$, $\models \delta_i \sqsubseteq \delta_{i+1}$.

Step 8. Now, we can define the inference relation $\vdash_{\mathcal{K}}$ as

$$C \vdash_{\mathcal{K}} D \text{ iff } \models C \sqcap \prod \tilde{\Phi} \sqcap \delta_i \sqsubseteq D,$$

where δ_i is the first $\{C\} \cup \tilde{\Phi}$ -consistent formula¹¹ of the sequence $\langle \delta_0, \dots, \delta_n \rangle$. This is the DL analogue as Step 6, Section 2.2.

Again, all steps require a decision procedure for the classical entailment relation \models of DLs. We can redefine the properties characterizing a rational consequence relation into the framework of DLs.

We can show that

Proposition 5.1. $\vdash_{\mathcal{K}}$ is a rational consequence relation containing \mathcal{K} .

That is, the analogous properties of the propositional rational consequence relation are satisfied, namely:

$$\begin{array}{l} \text{(REF)} \quad C \vdash_{\mathcal{K}} C \\ \text{(LLE)} \quad \frac{C \vdash_{\mathcal{K}} E \quad \models_{\mathcal{T}} C = D}{D \vdash_{\mathcal{K}} E} \qquad \text{(RW)} \quad \frac{C \vdash_{\mathcal{K}} D \quad \models_{\mathcal{T}} D \sqsubseteq E}{C \vdash_{\mathcal{K}} E} \\ \text{(CT)} \quad \frac{C \sqcap D \vdash_{\mathcal{K}} E \quad C \vdash_{\mathcal{K}} D}{C \vdash_{\mathcal{K}} E} \qquad \text{(CM)} \quad \frac{C \vdash_{\mathcal{K}} E \quad C \vdash_{\mathcal{K}} D}{C \sqcap D \vdash_{\mathcal{K}} E} \\ \text{(OR)} \quad \frac{C \vdash_{\mathcal{K}} E \quad D \vdash_{\mathcal{K}} E}{C \sqcup D \vdash_{\mathcal{K}} E} \qquad \text{(RM)} \quad \frac{C \vdash_{\mathcal{K}} D \quad C \not\vdash_{\mathcal{K}} \neg E}{C \sqcap E \vdash_{\mathcal{K}} D} \end{array}$$

11. That is, $\not\models C \sqcap \prod \tilde{\Phi} \sqsubseteq \neg \delta_i$.

Example 5.2. Consider example 5.1, where additionally we also add a role *Born* ('*Born*(x, y)' is read as ' x is born from y '), and a concept e (*Egg*). Consider

$$\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle ,$$

where

$$\mathcal{T} = \{p \sqsubseteq b, i \sqsubseteq \neg fi\}$$

$$\mathcal{D} = \{p \sqsubseteq \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top, b \sqsubseteq \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top, b \sqsubseteq \exists \text{Born}.e\} .$$




Figure 12: Example 5.2

Now (**Step 1**), we build the correspondent net $\mathcal{N}_{\mathcal{K}}$ (figure 12), and we obtain (**Step 2**) that

$$\mathcal{D}' = \mathcal{D} \cup \{p \sqsubseteq \exists \text{Born}.e\} .$$

Then we move to the rational closure. The pair $\langle \mathcal{T}, \mathcal{D}' \rangle$ is changed into (**Step 3**)

$$\begin{aligned} \mathcal{D}^* = \{ & p \sqcap \neg b \sqsubseteq \perp, i \sqcap fi \sqsubseteq \perp, \\ & p \sqsubseteq \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top, b \sqsubseteq \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top, \\ & b \sqsubseteq \exists \text{Born}.e, p \sqsubseteq \exists \text{Born}.e \} . \end{aligned}$$

The set of the materialisations of \mathcal{D}^* is (**Step 4**)

$$\begin{aligned} \overline{\mathcal{D}^*} &= \{p \wedge \neg b \supset \perp, i \sqcap fi \supset \perp, p \supset \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top, \\ & b \supset \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top, b \supset \exists \text{Born}.e, p \supset \exists \text{Born}.e\} \\ \mathfrak{A}_{\mathcal{D}^*} &= \{p \sqcap \neg b, i \sqcap fi, p, b\} . \end{aligned}$$

We obtain now (**Step 5**) the exceptionality ranking of the conditionals:

$$\begin{aligned} \mathcal{E}_0 &= \{p \sqcap \neg b \sqsubseteq \perp, i \sqcap fi \sqsubseteq \perp, p \sqsubseteq \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top, \\ & b \sqsubseteq \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top, b \sqsubseteq \exists \text{Born}.e, p \sqsubseteq \exists \text{Born}.e\} \\ \mathcal{E}_1 &= \{p \sqcap \neg b \sqsubseteq \perp, i \sqcap fi \sqsubseteq \perp, p \sqsubseteq \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top, p \sqsubseteq \exists \text{Born}.e\} \\ \mathcal{E}_2 &= \{p \sqcap \neg b \sqsubseteq \perp, i \sqcap fi \sqsubseteq \perp\} \\ \mathcal{E}_3 &= \{p \sqcap \neg b \sqsubseteq \perp, i \sqcap fi \sqsubseteq \perp\} = \mathcal{E}_2 . \end{aligned}$$

From this we get the ranking values of every conditional in \mathcal{D}^* : namely,

$$\begin{aligned} 0 &= \mathfrak{r}(b \sqsubseteq \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top) = \mathfrak{r}(b \sqsubseteq \exists \text{Born}.e) \\ 1 &= \mathfrak{r}(p \sqsubseteq \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top) = \mathfrak{r}(p \sqsubseteq \exists \text{Born}.e) \\ \infty &= \mathfrak{r}(p \sqcap \neg b \sqsubseteq \perp) = \mathfrak{r}(i \sqcap fi \sqsubseteq \perp) . \end{aligned}$$

From this ranking, we obtain (**Steps 6-7**) a background theory

$$\tilde{\mathcal{T}} = \{\top \sqsubseteq \neg(p \sqcap \neg b), \top \sqsubseteq \neg(i \sqcap fi)\} ,$$

and a default-assumption set $\tilde{\Delta} = \{\delta_0, \delta_1\}$, with

$$\begin{aligned} \delta_0 &= (b \supset \forall \text{Prey}.i \sqcap \exists \text{Prey}.\top) \sqcap (b \supset \exists \text{Born}.e) \sqcap \\ &\quad (p \supset \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top) \sqcap (p \supset \exists \text{Born}.e) \\ \delta_1 &= (p \supset \forall \text{Prey}.fi \sqcap \exists \text{Prey}.\top) \sqcap (p \supset \exists \text{Born}.e) \end{aligned}$$

to be used for the definition of the consequence relation $\vdash_{\mathcal{K}}$ (**Step 8**).

For example, we derive that typical penguins' preys are fishes, i.e.,

$$p \vdash_{\mathcal{K}} \forall \text{Prey}.fi ,$$

and not insects, i.e.,

$$p \vdash_{\mathcal{K}} \forall \text{Prey}.\neg i ,$$

and that also penguins are born from eggs, i.e.,

$$p \vdash_{\mathcal{K}} \exists \text{Born}.e ,$$

that would not be derivable from the rational closure, as presented by Casini and Straccia (2010).

Computational Complexity. From a computational complexity point of view, as deciding entailment in \mathcal{ALC} is ExpTime-complete (Donini & Massacci, 2000) and, as for Section 2.2, the number of entailment tests is polynomially bounded by the size of the knowledge base, following exactly the same procedures defined for the propositional case, we conclude that

Proposition 5.2. *Deciding $C \vdash_{\mathcal{K}} D$ in \mathcal{ALC} is an ExpTime-complete problem.*

5.1 Closure Operation over Individuals

So far we left out the ABox, but we are going to consider it here. The procedure for the ABox is built over the above procedure for the TBox, that is, we consider a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ such that all the strict knowledge has already been moved into \mathcal{T} , i.e., in \mathcal{D} we do not have axioms with ∞ as ranking value (that is, they correspond to the sets $\tilde{\mathcal{T}}$ and $\tilde{\mathcal{D}}$ obtained using the procedure in the previous section). The basic idea of the following procedure is to consider each individual named in the ABox as much typical as possible, that is, to associate to it all the possible defeasible information that is consistent with the rest of the knowledge base. In order to apply the defeasible information locally to each

individual, we encode such information using the materialisations of the inclusion axioms, i.e., the set $\Delta = \langle \delta_1, \dots, \delta_n \rangle$, s.t. $\models \delta_i \sqsubseteq \delta_{i+1}$ for $1 \leq i < n$ (the set $\tilde{\Delta}$ in the section above¹²). We want to be able to associate to each individual $\mathbf{a} \in \mathcal{O}$ (with \mathcal{O} being the set of the individuals named in the ABox) the strongest formula δ_i that is consistent with the knowledge base. In such a way we define a new knowledge base $\tilde{\mathcal{K}} = \langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$, that we call a *ABox extension* of the knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$.

Definition 5.1 (ABox extension). *Given a knowledge base $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, a knowledge base $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$ is an ABox extension of $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ iff*

- $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$ is classically consistent and $\mathcal{A} \subseteq \mathcal{A}_{\mathcal{D}}$.
- $\mathcal{A}_{\mathcal{D}} \setminus \mathcal{A}$ is composed of all the assertions $\mathbf{a}:C$ such that $\mathbf{a} \in \mathcal{O}$ and $C = \delta_i$ for some i , and for every δ_h s.t. $h < i$,

$$\langle \mathcal{T}, \mathcal{A}_{\mathcal{D}} \cup \{\mathbf{a}:\delta_h\} \rangle \models \perp$$

The above definition identifies the extensions of the original ABox \mathcal{A} s.t. to every individual is associated all the defeasible information that is consistent with the rest of the knowledge base. Still, the main problem is that, since the individuals can be related to each other through roles, the possibility of associating a default concept to an individual can be influenced by the default information associated to other individuals, as shown in the following example.

Example 5.3. *Consider $\mathcal{K} = \langle \mathcal{A}, \mathcal{D} \rangle$, with $\mathcal{A} = \{(\mathbf{a}, \mathbf{b}):R\}$ and $\mathcal{D} = \mathcal{D}_0 = \{\top \sqsubseteq A \sqcap \forall R. \neg A\}$ (hence we have $\Delta = \langle \delta_0 \rangle = \langle A \sqcap \forall R. \neg A \rangle$). If we associate δ_0 to \mathbf{a} , we obtain $\mathbf{b}:\neg A$ and we cannot associate δ_0 to \mathbf{b} ; on the other hand, if we apply δ_0 to \mathbf{b} , we derive $\mathbf{b}:A$ and we are not anymore able to associate δ_0 to \mathbf{a} . Hence, we define two possible rational extensions of \mathcal{K} . \square*

This implies that, given a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, even if the closure of $\langle \mathcal{T}, \mathcal{D} \rangle$ is always unique there is the possibility that we have more than one ABox extensions. A simple procedure to obtain all the possible extensions of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, with \mathcal{O} the set of the individuals in named in \mathcal{A} , is the following:

Definition 5.2 (Procedure for ABox extensions).

- Consider the set \mathfrak{S} of all the linear orders of the individuals in \mathcal{O} ;
- For each $s = \langle \mathbf{a}_1, \dots, \mathbf{a}_m \rangle$ in \mathfrak{S} do:
 - Set $j := 1$
 - Set $\mathcal{A}_{\mathcal{D}}^s := \mathcal{A}$
 - Repeat until $j = m + 1$:
 - * Find the first default $\delta_i \in \Delta$ such that $\langle \mathcal{A}_{\mathcal{D}}^s \cup \{\mathbf{a}_j:\delta_i\}, \mathcal{T} \rangle \not\models \top \sqsubseteq \perp$.
 - * $\mathcal{A}_{\mathcal{D}}^s := \mathcal{A}_{\mathcal{D}}^s \cup \{\mathbf{a}_j:\delta_i\}$.

12. Note that, given a conditional knowledge base, the procedure to determine $\tilde{\mathcal{T}}, \tilde{\mathcal{D}}$ and $\tilde{\Delta}$ must be done once for all.

- * $j = j + 1$
- return $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$

where we indicate by $\mathcal{A}_{\mathcal{D}}^s$ the ABox extension of \mathcal{A} obtained using the sequence s .

Now, it can be shown that

Proposition 5.3. *Given a linear order of the individuals in \mathcal{K} , the above procedure determines an ABox extension of \mathcal{K} . Vice-versa, every ABox extension of \mathcal{K} corresponds to the knowledge base generated by some linear order of the individuals in \mathcal{O} .*

For instance, related to Example 5.3, we obtain the extension $\tilde{\mathcal{K}}_1 = \langle \mathcal{A} \cup \{\mathbf{a}:A \sqcap \forall R.\neg A\}, \emptyset \rangle$ from the order $\langle \mathbf{a}, \mathbf{b} \rangle$, while $\tilde{\mathcal{K}}_2 = \langle \mathcal{A} \cup \{\mathbf{b}:A \sqcap \forall R.\neg A\}, \emptyset \rangle$ is obtained from the order $\langle \mathbf{b}, \mathbf{a} \rangle$.

Example 5.4. *Refer to Example 5.1, and let $\mathcal{K} = \{\mathcal{A}, \mathcal{T}, \mathcal{D}\}$, where*

$$\begin{aligned} \mathcal{A} &= \{\mathbf{a}:p, \mathbf{b}:b, (\mathbf{a}, \mathbf{c}):Prey, (\mathbf{b}, \mathbf{c}):Prey\} \\ \mathcal{T} &= \{p \sqsubseteq b, i \sqsubseteq \neg fi\} \\ \mathcal{D} &= \{b \sqsupseteq f, b \sqsupseteq \forall Prey.i, p \sqsupseteq \neg f, p \sqsupseteq \forall Prey.fi\} \end{aligned}$$

From such a knowledge base we define a set $\Delta = \{\delta_0, \delta_1\}$ with

$$\begin{aligned} \delta_0 &= (b \supset f) \sqcap (b \supset \forall Prey.i) \sqcap (p \supset \neg f) \sqcap (p \supset \forall Prey.fi) \\ \delta_1 &= (p \supset \neg f) \sqcap (p \supset \forall Prey.fi). \end{aligned}$$

If we consider an order where \mathbf{a} comes before \mathbf{b} , then we associate δ_1 to \mathbf{a} , and consequently \mathbf{c} is presumed to be a fish and we are prevented in the association of δ_0 to \mathbf{b} . If we consider \mathbf{b} before \mathbf{a} , \mathbf{c} is not a fish and we cannot apply δ_1 to \mathbf{a} .

Now, if we fix a priori a linear order s on the individuals, we can say that $\mathbf{a}:C$ is a *defeasible consequence* of \mathcal{K} w.r.t. the order s , written ' $\mathcal{K} \Vdash_s \mathbf{a}:C$ ', iff $\tilde{\mathcal{K}} \models \mathbf{a}:C$, where $\tilde{\mathcal{K}}$ is the ABox extension generated from \mathcal{K} based on the order s .

For instance, related to Example 5.3 and order $s_1 = \langle \mathbf{a}, \mathbf{b} \rangle$, we may infer that $\mathcal{K} \Vdash_{s_1} \mathbf{a}:A$, while with order $s_2 = \langle \mathbf{b}, \mathbf{a} \rangle$, we may infer that $\mathcal{K} \Vdash_{s_2} \mathbf{b}:A$.

The interesting point of such a consequence relation is that it still satisfies the properties of a *rational* consequence relation in the following way.

$$\begin{array}{l}
 (REF_{DL}) \quad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \text{ for every } \mathbf{a}:C \in \mathcal{A} \\
 (LLE_{DL}) \quad \frac{\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \models D = E}{\langle \mathcal{A} \cup \{\mathbf{b}:E\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C} \\
 (RW_{DL}) \quad \frac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \models C \sqsubseteq D}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:D} \\
 (CT_{DL}) \quad \frac{\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{b}:D}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C} \\
 (CM_{DL}) \quad \frac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{b}:D}{\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C} \\
 (OR_{DL}) \quad \frac{\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \langle \mathcal{A} \cup \{\mathbf{b}:E\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C}{\langle \mathcal{A} \cup \{\mathbf{b}:D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C} \\
 (RM_{DL}) \quad \frac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C \quad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash_s \mathbf{b}:\neg D}{\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C}
 \end{array}$$

We can show that

Proposition 5.4. *Given \mathcal{K} and a linear order s of the individuals in \mathcal{K} , the consequence relation \Vdash_s satisfies the properties (REF_{DL}) – (RM_{DL}).*

Note that from a computational complexity point of view, as entailment w.r.t. a \mathcal{ALC} TBox is ExpTime-complete (Donini & Massacci, 2000) and the number of individuals in \mathcal{K} is linearly bounded by $|K|$, we get immediately

Proposition 5.5. *Deciding $\mathcal{K} \Vdash_s \mathbf{a}:C$ in \mathcal{ALC} is a ExpTime-complete problem.*

In the presence of multiple ABox extensions, we can also define the inference relation \Vdash , a more conservative inference relation independent from any order on the individuals, that corresponds to the intersection of all the inference relations \Vdash^s modelling a rational extension.

$$\Vdash = \bigcap \{ \Vdash^s \mid s \text{ is a linear order on the elements of } \mathcal{O} \}$$

However, there is the possibility that we lose the property of rational monotonicity, as shown in the following example.

Example 5.5. *Consider the knowledge base $\langle \mathcal{A}, \mathcal{D} \rangle$ s.t. $\mathcal{A} = \{(\mathbf{a}, \mathbf{b}):R\}$ and $\mathcal{D} = \mathcal{D}_0 \cup \mathcal{D}_1$, with $\mathcal{D}_0 = \{\top \sqsubseteq A \sqcap \forall R. \neg A, \top \sqsubseteq B\}$ and $\mathcal{D}_1 = \{\neg A \sqsubseteq \neg B, \neg \forall R. \neg A \sqsubseteq B\}$ (where \mathcal{D}_0 and \mathcal{D}_1 are the sets of the conditionals of rank 0 and of the conditionals of rank 1, respectively). We can define two sequences on the individuals, $s = \langle \mathbf{a}, \mathbf{b} \rangle$ and $s' = \langle \mathbf{b}, \mathbf{a} \rangle$, each of them defining a different rational extension (\Vdash^s and $\Vdash^{s'}$), and let $\Vdash = \Vdash^s \cap \Vdash^{s'}$. We have that $\langle \mathcal{A}, \mathcal{D} \rangle \Vdash \mathbf{a}:B$, since in both the extensions $\mathbf{a}:B$ holds (in \Vdash^s because of the axiom $\top \sqsubseteq B$ and in $\Vdash^{s'}$ for the axiom $\neg \forall r. \neg A \sqsubseteq B$) while we have $\langle \mathcal{A}, \mathcal{D} \rangle \not\Vdash \mathbf{a}:A$, since $\langle \mathcal{A}, \mathcal{D} \rangle \not\Vdash^{s'} \mathbf{a}:A$. However, $\langle \mathcal{A} \cup \{\mathbf{a}:\neg A\}, \mathcal{D} \rangle \not\Vdash \mathbf{a}:B$, since $\langle \mathcal{A} \cup \{\mathbf{a}:\neg A\}, \mathcal{D} \rangle \not\Vdash^s \mathbf{a}:B$.*

We have no increase in the computational complexity of the decision procedure: assuming that the number of individuals named in the ABox is n , we have to perform a \Vdash^s -test for each possible sequences s defined on the n individuals. That is, in the worst case we need to do $n!$ \Vdash^s tests, each of which can be done in time $O(2^{|\mathcal{K}|^k})$ for some k . Now, it can be shown that¹³ $n! < 2^{n^2}$ and, thus, the decision problem for \Vdash remains in ExpTime.

Proposition 5.6. *Deciding $\mathcal{K} \Vdash \mathbf{a}:C$ in \mathcal{ALC} is a ExpTime-complete problem.*

Notwithstanding, we conjecture that in many (probably most) of the real-world cases, a knowledge base would have a single rational ABox extension, and in such cases (RM_{DL}) is still valid. To check whether a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ has a single rational ABox extension, it is sufficient to associate to each individual in \mathcal{O} the strongest δ_i modulo consistency w.r.t $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, exactly as in the procedure in Definition 5.2, but doing the consistency check of $\mathbf{a}_j:\delta_i$ w.r.t. the original \mathcal{A} instead that w.r.t. $\mathcal{A}_{\mathcal{D}}^s$. In the end, check whether $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$ is consistent; in such a case we have obtained the only rational ABox extension of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$.

The following is a knowledge base with a unique ABox extension.

Example 5.6. *Consider the KB in Example 5.4, where $(\mathbf{b}, \mathbf{c}):Prey$ is replaced with $(\mathbf{b}, \mathbf{d}):Prey$. Then, whatever is the order on the individuals, we obtain the following association between the default formulae and the individuals: $\mathbf{a}:\delta_1$, $\mathbf{b}:\delta_0$, $\mathbf{c}:\delta_0$, and $\mathbf{d}:\delta_0$. Using the information in these defaults, we obtain an unique default-assumption extension. \square*

A semantic characterization of \Vdash and \Vdash^s , making use of preferential DL models, is presented by Casini et al. (2013).

Now let's briefly consider some heuristics that are useful in case we want to present to the system specific ABox queries. Assume we want to know if a particular individual \mathbf{a} presumably falls under a concept C , and we want to draw the safest possible conclusion. In the presence of multiple acceptable extensions, the classical solution is to use a *skeptical* approach, *i.e.*, to use the inference relation \Vdash , corresponding to the intersection of all the inference relations associated to each possible ordering s of the individuals appearing in \mathcal{A} .

As we have seen above, in case of multiple rational extensions the computational complexity of the \Vdash decision problem does not rise w.r.t the classical \mathcal{ALC} decision problem. Moreover, in case of multiple extensions, the amount of defeasible information associable to an individual \mathbf{a} can be influenced only by the individuals related to it by means of a role: it is immediate to see that if there is no role-connection in the ABox between two individuals \mathbf{a} and \mathbf{b} , then the information that is associated to \mathbf{a} does not influence at all the amount of defeasible information that we can associate to \mathbf{b} , and the other way around. Hence, we can ease the decisions w.r.t. the ABox introducing the notion of *cluster*, *i.e.*, a set of individuals named in the ABox that are linked by means of a sequence of role connections. To do so, given an ABox \mathcal{A} , we indicate with Q the symmetric and transitive closure of all the roles in our vocabulary, *i.e.*, the symmetric and transitive closure of $\bigcup \mathcal{R}$.

Definition 5.3 (Cluster). *Define Q as the reflexive, symmetric and transitive closure of $\bigcup \mathcal{R}$. Given an individual $\mathbf{a} \in \mathcal{O}$, we call the cluster of \mathbf{a} the set $[\mathbf{a}]$ of the individuals*

13. This can be shown by induction on n , or see *e.g.* <http://lifecs.likai.org/2012/06/better-upper-bound-for-factorial.html>.

connected to \mathbf{a} through Q :

$$[\mathbf{a}] = \{\mathbf{b} \in \mathcal{O} \mid Q(\mathbf{a}, \mathbf{b})\} .$$

Hence, in order to know what we can presumably conclude about \mathbf{a} , it is sufficient to determine \Vdash^s w.r.t. each sequence s of individuals in $[\mathbf{a}]$. Let $\mathcal{A}_{[\mathbf{a}]}$ be the ABox obtained restricting \mathcal{A} to the statements containing individuals in $[\mathbf{a}]$; the query $\mathbf{a}:C$ is clearly decidable using only $\mathcal{A}_{[\mathbf{a}]}$.

Proposition 5.7. $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash \mathbf{a}:C$ iff $\langle \mathcal{A}_{[\mathbf{a}]}, \mathcal{T}, \mathcal{D} \rangle \Vdash^s \mathbf{a}:C$ for every ordering s of the individuals in $\mathcal{A}_{[\mathbf{a}]}$.

If we have a query about an individual \mathbf{a} s.t. \mathbf{a} is not named in the ABox ($\mathbf{a} \notin \mathcal{O}$), we do not have any constraints defined in the ABox about \mathbf{a} , *i.e.*, we only know $\mathbf{a}:\top$; hence, for each individual not appearing in the ABox, we can associate with it the strongest default concept consistent with \mathcal{T} , that is δ_0 : for any \mathbf{a} s.t. $\mathbf{a} \notin \mathcal{O}$, we can derive that presumably $\mathbf{a}:C$ holds iff $\langle \mathcal{A}_{\mathbf{a}}, \mathcal{T} \rangle \models \mathbf{a}:C$, where $\mathcal{A}_{\mathbf{a}} = \mathcal{A} \cup \{\mathbf{a}:\delta_0\}$.

6. Comparison with Related Work

Between non-monotonic logics, the so-called preferential approach can be distinguished from the other various proposals (as Reiter's defaults, modal approaches, defeasible inheritance...) mainly due its logical properties, since the former approach is committed to the satisfaction of some desirable structural properties of the consequence relation (see Section 2.2). On the other hand, considering them from the point of view of the inferential capacity the preferential approach often results weaker than other proposals, since often there are desirable, intuitive conclusions that we cannot derive (see Remark 1).

In our proposal we have tried to combine the classical rational closure with inheritance networks in order to overcome the inferential limits without prejudicing the logical properties of the consequence relation.

In Section 3 we also present an alternative way to reason about defeasible inheritance. Despite our proposal has been presented mainly to integrate it with propositional language and rational closure, it results an interesting approach per se, and in Appendix A we compare it with Horty's classical skeptical extension (Horty, 1994, sect. 2-3) and Sandewall's landmark examples (Sandewall, 2010).

As indicated in the introduction, there have been many papers aimed at the implementation of non-monotonic reasoning into the DL formalisms. For most of such proposals the comparison with our approach has to be done considering the different non-monotonic formalisms, independently from the DL-environment. So we refer to Makinson's work (1994) for a comparison between the various non-monotonic approaches.

In the last years the main proposals for the implementation of nonmonotonic reasoning in DLs have been connected to two approaches: the preferential one and circumscription.

In the preferential approach, the work by Britz and al. (Britz et al., 2008; Britz, Meyer, & Varzinczak, 2011) on preferential DL semantics is strongly connected to our approach, and one of the results has been the semantic characterization of rational closure cited above (Britz et al., 2013).

Still very close to our approach is the work by Giordano et al. (2012b), that is based too on a preferential approach. The conclusions that we can derive using the logic $\mathcal{ALC} + T_{min}$

are intuitive, but the complexity of the decision problem for the ABox is $\text{co-NExp}^{\text{NP}}$ (Giordano et al., 2012b, Thm. 13), and the procedure cannot be reduced to classical entailment.

Among the proposals based on circumscription, the work by Bonatti et al. (2009) is particularly representative. From the point of view of the quality of the inferences, in such a proposal it results more difficult w.r.t. the preferential approach to draw the expected conclusions. For example, assume that our knowledge base contains the information that mammals typically live on land, but that whales are abnormal mammals that do not live on the land, and the ABox contains the information $\mathbf{a}:Mammal \sqcap \neg Whale$. Not knowing anything else about the individual \mathbf{a} , we would like our reasoning system to reason on the assumption that we are dealing with a typical mammal (since, moreover, it is specified that \mathbf{a} is not a whale) and hence being able to derive that \mathbf{a} lives on the land. However, using circumscription, the conclusions we can draw changes w.r.t. which concepts the user decides to keep *fixed* or *varying* (a non-trivial choice), and the results can be that we are not able to derive $\mathbf{a}:\exists Habitat.Land$, that we are able to derive it, or we can even derive that whales do not exist (Bonatti et al., 2009, sect. 2.1). In our proposal instead, we can formalize the problem with a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ with $\mathcal{A} = \{\mathbf{a}:Mammal\}$ (we do not need to specify that it is not a whale), $\mathcal{T} = \{Whale \sqsubseteq Mammal, Whale \sqsubseteq \neg \exists Habitat.Land\}$ and $\mathcal{D} = \{Mammal \sqsupseteq \exists Habitat.Land\}$; without needing any kind of choice from the user, the system can derive automatically $\mathbf{a}:\exists Habitat.Land$. Moreover, we have seen that the computational cost of our procedure involving an ABox is exponential, while in the circumscription case, for languages analogous to \mathcal{ALC} , the complexity of the instance problem is $\text{co-NExp}^{\text{NP}}$ (Bonatti et al., 2009, sect. 4.1.1). Some of the issues just discussed have been addressed and solved by Bonatti et al. (2010, 2011a), but only for circumscriptive systems that are specifically built for low-complexity DLs such as \mathcal{EL} .

On the other hand, procedures based on circumscription are able to derive defeasible information about individuals that are implicit in the ABox, that is, they can, for example, conclude from $\mathbf{a}:\exists R.C$ and $C \sqsupseteq D$ that presumably $\mathbf{a}:\exists R.D$. Our procedure involving ABoxes is not able to make such a kind of derivation, since we can add defeasible information only to the individuals named in the ABox. We are working on a further refinement of the procedure in order to deal also with the implicit individuals; a first attempt to take under consideration also such individuals has been proposed in a previous publications (Casini & Straccia, 2010, pp. 9-10), adding a completion procedure for the ABox in order to explicitly name in the ABox the implicit individuals, but such a procedure needs to be refined.

Among the proposals regarding the introduction of probabilistic reasoning in DLs, Lukasiewicz (2008) presents a combination of nonmonotonic and probabilistic reasoning. The nonmonotonic part is based on the preferential approach, and he too presents a construction that augments the inferential power of rational closure. Let us consider his proposal by eliminating the probabilistic part, *i.e.*, considering only conditionals associated to probability 1 (conditionals $(\psi|\phi)[1, 1]$ in his notation), since they convey the same meaning as our defeasible conditionals. His procedure seems to give back the same results as ours in most of the cases, but the two proposals differ in the general approach: he proceeds with a refinement of the ranking structure, while we use a refinement of the content of the knowledge base. In fact, the behaviour of the two can differ. Consider for example a knowledge base $\langle \emptyset, \mathcal{D} \rangle$, with $\mathcal{D} = \{a \sqsupseteq b, a \sqsupseteq d, b \sqsupseteq c, c \sqsupseteq e, d \sqsupseteq \neg e\}$ (in Lukasiewicz's notation it would correspond to a knowledge base with an empty TBox and a set P containing the conditionals

$(b|a)[1, 1], \dots)$.¹⁴ In our construction we can derive $a \sim c$, since there is simply a duct $\langle a, b, c \rangle$ from a to c , while in Łukasiewicz’s approach we cannot, since we have to consider what follows from all the three preferred subset of the knowledge base that are consistent with a (that are $\{a \sqsubseteq b, a \sqsubseteq d, b \sqsubseteq c, c \sqsubseteq e\}$, $\{a \sqsubseteq b, a \sqsubseteq d, b \sqsubseteq c, d \sqsubseteq \neg e\}$, $\{a \sqsubseteq b, a \sqsubseteq d, c \sqsubseteq e, d \sqsubseteq \neg e\}$), and $a \sim c$ does not follow from the latter.

7. Conclusions

By combining the classical rational closure with ideas from defeasible inheritance networks, we have proposed a new rational consequence relation that overcomes some limits of both the formalisms. By doing so, we have extended the defeasible inference capabilities by allowing an atypical class still to inherit some properties from its superclass while maintaining the desired logical properties of rational closure. The table below summarizes the structural properties satisfied by the systems taken under consideration:¹⁵

	Horty	IN	BIN	PL	DL
REF	•	•	•	•	•
CT		•	•	•	•
CM		•	•	•	•
LE	•	•	•	•	•
RW	•	•	•	•	•
OR			•	•	•
RM				•	•

As we can see, our proposals for defeasible inheritance-based propositional logic and Description Logics still satisfy all axioms of classical rational closure. Another feature is that our method does uniquely require the existence of a decision procedure of classical entailment and, thus, can be implemented on top of exiting propositional SAT solvers and DL reasoners. Since we have introduced a procedure that is interesting also from the point of view of inheritance nets, we have presented the more comprehensive procedure for logical knowledge bases making use of the nets’ formalism; notwithstanding, we are confident that the procedure can be reformulated avoiding the shift from one formalism to another.

The procedure presented for \mathcal{ALC} can be straightly extended to languages more expressive than \mathcal{ALC} ; on the other hand, at present the procedure needs a language that is closed under propositional connectives, and hence we need to augment the expressivity of the language in order to apply it to less expressive DL languages as \mathcal{EL} , forcing in this way an increase in the computational costs w.r.t. the classical decision problem.

Hence, looking for an adaptation of the procedure appropriate for dealing with tractable, less expressive DLs is one of the main open problems of the present proposal, together with a more proper semantic characterization of the procedure and the ability of reasoning in a defeasible way about implicit individuals, as discussed in the previous section.

14. The example corresponds to the structure of Example.A.4 in Appendix A, just eliminating the link at the center of the figure.

15. IN, BIN, PL, DL stand for our proposals for INs, Boolean INs, propositional logic and DLs, respectively.

Appendix A. Examples

We are going to validate our decision procedures for inheritance nets against some significant examples proposed by Horty (1994) and Sandewall (2010). We shall use BINs, but, in order to simplify the graphical representation, we shall also use \nrightarrow as a macro, as explained in Section 3.2. As a side effect, we obtain an analogue behaviour for the propositional and DL cases too.

Example A.1 (Horty, 1994, ex. Γ_{12}). Consider the net in Figure 13. Horty claims that the desirable conclusion should be $a \sim p$, since “In the environment of mixed nets [...] certain kinds of compound defeasible paths can legitimately be thought to carry immediate information - namely, those paths consisting of a single defeasible link followed by a strict end segment, of any length” (Horty, 1994, p. 143). That corresponds to the condition of (RW), that, as we have seen in Section 3.1.2, our procedure satisfies.




Figure 13: Example A.1

Indeed, we translate the net into the KB $\mathcal{K} = \langle \Phi, \Delta \rangle$, with $\Phi = \{a \supset s, s \supset t, r \supset p\}$ and $\Delta = \{s \supset r, t \supset \neg p\}$. We have $\Delta_{a,p} = \Delta$, $\Delta_{a,p} \models_{\Phi} \neg s$, and $\Delta_{a,p} \models_{\Phi} \neg a$. So the only implication in $\Delta_{a,p}$ with ranking greater than 0 is $s \supset r$, and, since $s \supset r \models_{\Phi} a \supset p$, we have $a \sim_{\mathcal{K}} p$.

Example A.2 (Horty, 1994, ex. Γ_{18}). Consider the net in Figure 14. In this example Horty claims that we should not prefer the negative path $\langle a, p, s \rangle$ over the positive path $\langle a, r, s \rangle$, since the negative link $a \nrightarrow q$ nullifies the path from p to r . Our procedure satisfies such a claim.




Figure 14: Example A.2

Indeed, we translate the net into the KB $\mathcal{K} = \langle \Phi, \Delta \rangle$, with $\Phi = \emptyset$ and $\Delta = \{a \supset p, a \supset r, p \supset q, q \supset r, r \supset s, a \supset \neg q, p \supset \neg s\}$. We have $\Delta_{a,s} = \Delta$, $\Delta_{a,s} \models \neg p$, and $\Delta_{a,p} \models \neg a$.

Therefore, we have $\mathcal{E}_1 = \{a \supset p, a \supset r, a \supset \neg q, p \supset q, p \supset \neg s\}$. Now, $\mathcal{E}_1 \models \neg a$, and a is the only node of rank 2, with $\mathcal{E}_2 = \{a \supset p, a \supset r, a \supset \neg q\}$. So we conclude neither $a \vdash s$ nor $a \vdash \neg s$.

Example A.3 (Sandewall, 2010, ex. B.1, Double Diamond). Consider the net $\mathcal{N} = \{\mathcal{S}, \mathcal{U}\}$ defined as (see Figure 15)

$$\begin{aligned} \mathcal{S} &= \{s \Leftrightarrow \neg \bar{s}, r \Leftrightarrow \neg \bar{r}\} \\ \mathcal{U} &= \{a \rightarrow t, a \rightarrow p, t \rightarrow s, p \rightarrow \bar{s}, s \rightarrow r, p \rightarrow q, q \rightarrow \bar{r}\} \end{aligned}$$




Figure 15: Double diamond.

The net \mathcal{N} is translated in the knowledge base $\mathcal{K} = \{\Phi, \Delta\}$ with $\Phi = \{\emptyset\}$ and $\Delta = \{a \supset t, a \supset p, t \supset s, p \supset \neg s, s \supset r, p \supset q, q \supset \neg r\}$.

Using our method, from this net we derive neither $a \vdash r$, nor $a \vdash \neg r$, as it should be. However, we can derive $a \vdash q$, which is a desirable result that is not derivable with Sandewall's approach.

Example A.4 (Sandewall, 2010, ex. B.2, Simonet's Scenario). Consider the net $\mathcal{N} = \{\mathcal{S}, \mathcal{U}\}$ defined as (Figure 16)

$$\begin{aligned} \mathcal{S} &= \{e \Leftrightarrow \neg \bar{e}\} \\ \mathcal{U} &= \{a \rightarrow b, a \rightarrow d, b \rightarrow c, c \rightarrow e, d \rightarrow c, d \rightarrow \bar{e}\} . \end{aligned}$$




Figure 16: Simonet's scenario.

The net \mathcal{N} is translated in the knowledge base $\mathcal{K} = \{\Phi, \Delta\}$ with $\Phi = \{\emptyset\}$ and $\Delta = \{a \supset b, a \supset d, b \supset c, c \supset e, d \supset c, d \supset \neg e\}$.

Using our method, from this net we derive $a \vdash \neg e$, as it should be, while Sandewall cannot derive it. Moreover, he derives neither $a \vdash c$ nor $b \vdash e$.

Example A.5 (Sandewall, 2010, ex. B.3, On-Path vs. Off-Path Preemption). Consider the net $\mathcal{N} = \{\mathcal{S}, \mathcal{U}\}$ defined as (see Figure 17)

$$\begin{aligned} \mathcal{S} &= \{wa \not\Leftarrow ga\} \\ \mathcal{U} &= \{c \rightarrow re, c \rightarrow ce, re \rightarrow e, ce \rightarrow e, e \rightarrow ga, re \rightarrow wa\} . \end{aligned}$$




Figure 17: On-Path vs Off-Path.

The net \mathcal{N} is translated in the knowledge base $\mathcal{K} = \{\Phi, \Delta\}$ with $\Phi = \{wa \supset \neg ga\}$ and $\Delta = \{c \supset re, c \supset ce, re \supset e, ce \supset e, e \supset ga, re \supset wa\}$.

The connection $c \rightarrow wa$ is valid only if we use a form of off-path preemption, while with on-path preemptions it is not derivable. In our setting, we derive $c \vdash wa$: even if we cannot consider our method a form of preemption (at least explicitly), as said in Section 3.1.5 our method gives back results that are analogous to those of off-path preemption.

The following examples use also conjunctions, so we are going to use BINs.

Example A.6 (Sandewall, 2010, ex. B.4, Juvenile Offender). The original “Juvenile offender” example is represented by the following net (see Figure 18) $\mathcal{N} = \{\mathcal{S}, \mathcal{U}\}$, with

$$\begin{aligned} \mathcal{S} &= \{b \Rightarrow g, b \Rightarrow m, p \Leftrightarrow \bar{p}, m, g \Leftrightarrow m \otimes g\} \\ \mathcal{U} &= \{g \rightarrow p, m \rightarrow \bar{p}\} , \end{aligned}$$

where p is read as ‘to be punished’, g as ‘guilty’, m as ‘is minor’, b is ‘Billy’. We want to express that m has the priority on g , and so that $b \rightarrow_{\mathcal{N}} \bar{w}$. Since in our nets it is possible to express conjunctions, we can solve the problem adding the link $m \otimes g \rightarrow \bar{p}$ and $m, g \Leftrightarrow m \otimes g$.

The net is translated in the knowledge base $\mathcal{K} = \{\Phi, \Delta\}$ with $\Phi = \{b \supset g, b \supset m\}$ and $\Delta = \{g \supset p, m \supset \neg p, m \wedge g \supset \neg p\}$. From this knowledge base we derive $b \vdash \neg p$, as expected.

Example A.7 (Sandewall, 2010, ex. B.5, Campus Residence Scenario). As in example A.6, in the “Campus residence” example there is the necessity to include priorities. We can solve




Figure 18: Juvenile offender.

the problem by inserting conjunctions. Indeed, the net $\mathcal{N} = \{\mathcal{S}, \mathcal{U}\}$ is composed by (see Figure 19)

$$\begin{aligned} \mathcal{S} &= \{t \Rightarrow ma, t \Rightarrow e, t \Rightarrow m, w \not\Leftarrow n, ma, e \Leftrightarrow^{\wedge} ma \otimes e\} \\ \mathcal{U} &= \{ma \otimes e \rightarrow w, m \rightarrow n\}, \end{aligned}$$

where ma is read as ‘is married’, e as ‘employed’, w as ‘lives in west apartments’, m as ‘male’, n as ‘lives in northern apartments’, t is ‘Tom’. We want to express that $ma \otimes e$ has the priority on m , and so that $t \rightarrow w$ results valid. to do so, we add the links $m \otimes ma \otimes e \rightarrow w$, and $m, ma \otimes e \Leftrightarrow^{\wedge} m \otimes ma \otimes e$.




Figure 19: Campus residence scenario.

The net is translated in the knowledge base $\mathcal{K} = \{\Phi, \Delta\}$ with $\Phi = \{t \supset ma, t \supset e, t \supset m, w \supset \neg n\}$ and $\Delta = \{ma \wedge e \supset w, m \supset n, m \wedge ma \wedge e \supset w\}$. From this knowledge base we derive $t \vdash w$, as expected.

We treat the last example as a propositional problem. Just to have a simpler net from the graphical point of view we do not put the negation nodes that are useless for the resolution, and we use a three-place conjunction link, just as a macro for the construction of the conjunction of three nodes.

Example A.8 (Sandewall, 2010, ex. B.6, Good Math Student Scenario). *The knowledge base $\mathcal{K} = \{\mathcal{T}, \mathcal{D}\}$ is composed by (see Figure 20)*

$$\begin{aligned} \mathcal{T} &= \{t \vdash gs, t \vdash fm, t \vdash mb, aa \vdash \neg ag\} \\ \mathcal{D} &= \{mb \sim ag, fm \wedge ag \sim mm, mm \wedge mb \wedge gs \sim aa, fm \sim pm\}, \end{aligned}$$




Figure 20: Good math student.

Sandewall identifies the following candidates for valid conclusions:

$$\begin{aligned} t \vdash ag \quad t \vdash aa \quad t \sim fm \wedge ag \\ t \vdash mm \quad t \vdash mm \wedge mb \quad t \vdash mm \wedge mb \wedge gs \quad t \vdash pm. \end{aligned}$$

We cannot conclude any of the conditionals above. The point is that the net is highly interconnected, especially by strict links, so that all the nodes have ranking 0, except for t and $mm \otimes mb \otimes gs$ that have ranking 1. Therefore, any of the conclusions above, with t as premise, cannot be obtained.

Appendix B. Proofs

Proposition 3.1. *Consider a net $\mathcal{N} = \langle \mathcal{S}, \emptyset \rangle$ and translate it into a set of propositional implications Φ . The following properties hold:*

1. *If \mathcal{N} is a consistent net, there is a valid strict positive (resp., negative) path $\langle p, \sigma, q \rangle$ from p to q , that is $\mathcal{N} \rightsquigarrow p \Rightarrow q$ (resp., $\mathcal{N} \rightsquigarrow p \not\Rightarrow q$), iff $\Phi \models p \supset q$ (resp., $\Phi \models p \supset \neg q$).*
2. *\mathcal{N} is inconsistent iff $\Phi \models \neg p$ for some $p \in \mathfrak{A}_\Phi$.*

3. *Deciding strict consequence can be done in polynomial time.*

Proof. First, it is easy to prove by induction on the length of the paths that, if there is a positive (resp., negative) path in \mathcal{S} from p to q , then $\Phi \vDash p \supset q$ (resp., $\Phi \vDash p \supset \neg q$). Hence we have:

- 1a) If \mathcal{N} is a consistent net and there is a valid strict positive (negative) path $\pi(p, \sigma, q)$ from p to q , then $\Phi \vDash p \supset q$ ($\Phi \vDash p \supset \neg q$).

Moreover, if \mathcal{N} is inconsistent, then we can conclude, for some p, q in it, $p \Rightarrow q$ and $p \not\Leftarrow q$, and so we have $\Phi \vDash p \supset q$ and $\Phi \vDash p \supset \neg q$, that is, $\Phi \vDash \neg p$. So we have one half of the second statement.

- 2a) If $\Phi \not\vdash \neg p$ for every $p \in \mathfrak{A}_\ell$, then \mathcal{N} is consistent.

Now we move to show the other halves of the statements.

- 1b) If \mathcal{N} is a consistent net and $\Phi \vDash p \supset q$ (resp., $\Phi \vDash p \supset \neg q$), then there is a valid strict positive (resp., negative) path $\pi(p, \sigma, q)$ from p to q .

In order to model the classical consequence relation \vDash , we use the propositional resolution method. We use the symbol \vdash_R to indicate the inferences in the resolution method.

Every element of Φ has the forms $p \supset q$ or $p \supset \neg q$. Such implications correspond, in the clause form, respectively to the clauses (*i.e.*, sets of literals) $\{\neg p, q\}$ and $\{\neg p, \neg q\}$. Call $\overline{\Phi}$ the set of the clauses corresponding to Φ . For example, if we assume that $\overline{\Phi} \vdash_R \{\neg p, q\}$ (that is, $\Phi \vDash p \supset q$), we have that the set of clauses $\overline{\Phi}, \{p\}, \{\neg q\}$ resolves into the empty set (that is, $\Phi, \neg(p \supset q) \vDash \perp$).

1. Positive case ($\mathcal{N} \rightsquigarrow p \Rightarrow q$): Assume $\Phi \vDash p \supset q$ for some $p, q \in P_{\mathcal{N}}$. In order to apply the refutational approach, $p \supset q$ has to be negated, and, since $\neg(p \supset q)$ is equivalent to $p \wedge \neg q$, we introduce the clauses $\{p\}, \{\neg q\}$. So, in the resolution approach $\Phi \vDash p \supset q$ is translated into $\overline{\Phi}, \{p\}, \{\neg q\} \vdash_R \emptyset$. We have also assumed that $\overline{\Phi}$ is consistent, and hence we have that $\overline{\Phi} \not\vdash_R \emptyset$. Since the set $\overline{\Phi}$ is composed only of pairs of literals (of form $\{\neg p, q\}$ or $\{\neg p, \neg q\}$), every reduction step between them gives back again a pair of literals (for example, from $\{\neg p, q\}$ and $\{\neg q, \neg r\}$, we obtain $\{\neg p, \neg r\}$). Therefore, in order to obtain the empty set we need to use both $\{p\}$ and $\{\neg q\}$ in the refutation procedure. So, the clause $\{\neg q\}$ must be eliminated using a clause containing q . Such a clause must necessarily have the form $\{\neg r_i, q\}$ for some r_i , and we obtain a new clause $\{\neg r_i\}$. Again, $\{\neg r_i\}$ can be eliminated only by a clause of form $\{\neg r_j, r_i\}$, obtaining a clause $\{\neg r_j\}$. Since the set Φ is finite, such a procedure has to terminate, and it can only be done if we obtain a clause $\{\neg p\}$, that can be resolved by $\{p\}$. That is, there has to be in Φ a clause $\{\neg p, r_l\}$, for some r_l , s.t. $\{\neg r_l\}$ appears in the reduction procedure. Now, the clauses of $\overline{\Phi}$ used in the reduction process correspond to a chain of links in the net \mathcal{S} : $p \Rightarrow r_l, \dots, r_i \Rightarrow q$, that define a valid path (p, σ, q) in \mathcal{S} .
2. Negative case ($\mathcal{N} \rightsquigarrow p \not\Leftarrow q$): assume $\Phi \vDash p \supset \neg q$ for some $p, q \in P_{\mathcal{S}}$. $p \supset \neg q$ has to be negated and it is translated into the clauses $\{p\}, \{q\}$. $\{q\}$ can be combined

with a clause of form $\{\neg q, \neg r_i\}$ (that represents $q \supset \neg r_i$, that, in turn, represents a link $q \not\Leftarrow r_i$) (case 1), or with a clause of form $\{\neg q, r_i\}$ (i.e., $q \supset r_i$, that, in turn, is $q \Rightarrow r_i$) (case 2).

- Case 1. We obtain the clause $\{\neg r_i\}$, and the procedure is the same of the positive case. The clauses used in the reduction process correspond to a strict negative path of form: $p \Rightarrow r_l, \dots, r_j \Rightarrow r_i, r_i \not\Leftarrow q$.
- Case 2. We obtain the clause $\{r_i\}$, and we would have combine it with a clause of form $\{\neg r_i, \neg r_j\}$, or of form $\{\neg r_i, r_j\}$. In the former case we move to case 1, ending with a strict negative path $p \Rightarrow r_h, \dots, r_l \Rightarrow r_j, r_j \not\Leftarrow r_i, r_i \Leftarrow q$. In the latter case we have a new clause $\{\neg r_j\}$, and we are again in case 2; however, such a procedure has to terminate, and, in order to terminate, such a reduction procedure has at some point to move into the case 1. The clauses used in the reduction procedure correspond to a path of form $p \Rightarrow r_i, \dots, r_j \Rightarrow r_k, r_k \not\Leftarrow r_l, r_l \Leftarrow r_m, \dots, r_n \Leftarrow q$, where the link $\not\Leftarrow$ corresponds to the shift to the case 1.

- 2b) Automatically, we have also that if \mathcal{N} is consistent, then $\Phi \not\models \neg p$ for every $p \in \mathfrak{A}_\ell$. Otherwise, we would have $\Phi \models p \supset q$ and $\Phi \models p \supset \neg q$ for some nodes p and q , that would imply, by the procedure above, the inconsistency of \mathcal{N} .
- 3) About the third point, consider that strict links can be encoded as 2-CNF formulae, also called Krom formulae, and the propositional 2-SAT problem is in P .

□

Proposition 3.2. *Consider a net \mathcal{N} . For every connection $p \sim_{\mathcal{N}} q$ (resp., $p \sim_{\mathcal{N}} \neg q$) validated by our procedure, there is a corresponding positive (resp., negative) potential path from p to q in the net \mathcal{N} .*

Proof. Define the set $\widehat{\Delta}_{p,q}$. Our procedure states that $p \sim_{\mathcal{N}} q$ iff $\widehat{\Delta}_{p,q} \cup \Phi \models p \supset q$ (and $p \sim_{\mathcal{N}} \neg q$ iff $\widehat{\Delta}_{p,q} \cup \Phi \models p \supset \neg q$). Following a procedure analogous to the one of the proof of proposition 3.1, we can show that the derivation of such implications is connected to the presence of positive (negative) potential paths in the net. □

Proposition 3.3. *\rightsquigarrow' satisfies (REF), (CT) and (CM).*

To prove this proposition we need first to prove the following lemma.

Lemma B.1. *Consider a net $\mathcal{N} = \langle \mathcal{S}, \mathcal{U} \rangle$ s.t. $\mathcal{N} \rightsquigarrow' p \star q$ ($\star \in \{\rightarrow, \not\rightarrow\}$). Call Φ the set of material implications corresponding to the links in \mathcal{S} , and \models_Φ the consequence relation obtained adding to \models the formulae in Φ as extra-axioms. Consider the net \mathcal{N}' obtained adding the link $p \star q$ in \mathcal{N} . Then, for every pair of nodes r, s in \mathcal{N} , we have that $\widehat{\Delta}_{r,s}^{\mathcal{N}}$ and $\widehat{\Delta}_{r,s}^{\mathcal{N}'}$ are equivalent w.r.t. \models_Φ .*

Proof. The two nets contain the same nodes. Given two nodes r and s , we have two possible cases: $p \star q \notin \mathcal{C}_{r,s}^{\mathcal{N}'}$ or $p \star q \in \mathcal{C}_{r,s}^{\mathcal{N}'}$. If $p \star q \notin \mathcal{C}_{r,s}^{\mathcal{N}'}$, we have $\mathcal{C}_{r,s}^{\mathcal{N}'} = \mathcal{C}_{r,s}^{\mathcal{N}}$, and consequently $\widehat{\Delta}_{r,s}^{\mathcal{N}'} = \widehat{\Delta}_{r,s}^{\mathcal{N}}$. If $p \star q \in \mathcal{C}_{r,s}^{\mathcal{N}'}$, we have that $\mathcal{C}_{r,s}^{\mathcal{N}'} \subseteq \mathcal{C}_{r,s}^{\mathcal{N}}$, and consequently $\mathcal{C}_{p,q}^{\mathcal{N}} \subseteq \mathcal{C}_{r,s}^{\mathcal{N}}$, and the

correspondent sets of courses in the two nets ($\mathcal{C}_{p,q}^{\mathcal{N}'}$ and $\mathcal{C}_{p,q}^{\mathcal{N}}$ on one hand and $\mathcal{C}_{r,s}^{\mathcal{N}'}$ and $\mathcal{C}_{r,s}^{\mathcal{N}}$ on the other) contain exactly the same nodes and the same links, apart, possibly, from $p \star q$. That corresponds to saying that $\Delta_{p,q}^{\mathcal{N}}$ and $\Delta_{p,q}^{\mathcal{N}'}$ on one hand, and $\Delta_{r,s}^{\mathcal{N}}$ and $\Delta_{r,s}^{\mathcal{N}'}$ on the other contain exactly the same implications, apart, possibly, from $p \supset l_q$ (where $l_q \in \{q, \neg q\}$), and exactly the same set of antecedents.

Since $\mathcal{N} \rightsquigarrow' p \star q$, we have $\hat{\Delta}_{p,q}^{\mathcal{N}} \models_{\Phi} p \supset l_q$, and consequently $\Delta_{p,q}^{\mathcal{N}} \models_{\Phi} p \supset l_q$. So, $\Delta_{p,q}^{\mathcal{N}}$ and $\Delta_{p,q}^{\mathcal{N}'}$ are \models_{Φ} -equivalent, and then also $\Delta_{r,s}^{\mathcal{N}}$ and $\Delta_{r,s}^{\mathcal{N}'}$ are \models_{Φ} -equivalent. So, since $\Delta_{r,s}^{\mathcal{N}}$ and $\Delta_{r,s}^{\mathcal{N}'}$ are \models_{Φ} -equivalent and contain the same set of antecedents, they generate the same ranking, that is, $\hat{\Delta}_{r,s}^{\mathcal{N}}$ and $\hat{\Delta}_{r,s}^{\mathcal{N}'}$ are equivalent w.r.t. \models_{Φ} . □

Now we can prove Proposition 3.3.

Proof. The satisfaction of (REF) is trivial, since every direct link is trivially valid.

(CT): assume $\mathcal{N} \rightsquigarrow' p \star q$ and $\mathcal{N}, p \star q \rightsquigarrow' r \star s$. We call \mathcal{N}' the net obtained adding $p \star q$ in \mathcal{N} . $\mathcal{N}' \rightsquigarrow' r \star s$ means that $\hat{\Delta}_{r,s}^{\mathcal{N}'} \models_{\Phi} r \supset l_s$ ($l_s \in \{s, \neg s\}$). By the lemma B.1 we have that $\hat{\Delta}_{r,s}^{\mathcal{N}'}$ and $\hat{\Delta}_{r,s}^{\mathcal{N}}$ are equivalent w.r.t. \models_{Φ} , so we have also $\hat{\Delta}_{r,s}^{\mathcal{N}} \models_{\Phi} r \supset l_s$, that is, $\mathcal{N} \rightsquigarrow' r \star s$.

(CM): assume $\mathcal{N} \rightsquigarrow' p \star q$ and $\mathcal{N} \rightsquigarrow' r \star s$. We call \mathcal{N}' the net obtained adding $p \star q$ in \mathcal{N} . $\mathcal{N} \rightsquigarrow' r \star s$ means that $\hat{\Delta}_{r,s}^{\mathcal{N}} \models_{\Phi} r \supset l_s$. By the lemma B.1 we have that $\hat{\Delta}_{r,s}^{\mathcal{N}}$ and $\hat{\Delta}_{r,s}^{\mathcal{N}'}$ are equivalent w.r.t. \models_{Φ} , so we have also $\hat{\Delta}_{r,s}^{\mathcal{N}'} \models_{\Phi} r \supset l_s$, that is, $\mathcal{N}, p \star q \rightsquigarrow' r \star s$. □

Proposition 3.5. \rightsquigarrow' satisfies (LE), (RW), and (Sup).

Proof. The proofs are immediate, considering our procedure and the fact that the implications corresponding to strict links are always present in each decision procedure. □

Proposition 3.6 A net \mathcal{N} is consistent iff we do not have a node p with $\tau(p) = \infty$, that is, we do not conclude both $\mathcal{N} \rightsquigarrow' p \rightarrow q$ and $\mathcal{N} \rightsquigarrow' p \not\rightarrow q$ for any pair p, q .

It is proved combining the following two lemmas.

Lemma B.2. Given a net \mathcal{N} , and two nodes p and q in it, we can conclude both $p \vdash_{\mathcal{N}} q$ and $p \vdash_{\mathcal{N}} \neg q$ iff $\tau(p) = \infty$.

Proof. From left to right. It is immediate from the definition of the ranking procedure: if $\tau(p) \neq \infty$, then the set of implications with at least the same ranking as p do not imply $\neg p$, and so they cannot imply both $p \supset q$ and $p \supset \neg q$ for any q .

From right to left. If $\tau(p) = \infty$, we have a set \mathcal{E}_{∞} ¹⁶ of implications (with p between their antecedents) such that $\mathcal{E}_{\infty} \models_{\Phi} \neg p$. That implies $\mathcal{E}_{\infty} \models_{\Phi} p \supset q$ and $\mathcal{E}_{\infty} \models_{\Phi} p \supset \neg q$ for whichever q , that is, $p \vdash_{\mathcal{N}} q$ and $p \vdash_{\mathcal{N}} \neg q$. □

Lemma B.3. Every node with ranking ∞ in a set of courses C has the same ranking value in every set of courses extending C .

Proof. It is immediate from the monotony of the consequence relations $\models_{\Phi, \mathcal{N}}$. □

16. \mathcal{E}_{∞} is the set of ∞ -ranked defaults.

From these lemmas we conclude that if a node p has infinite ranking in a net, then the net is inconsistent ($p \sim_{\mathcal{N}} q$ and $p \sim_{\mathcal{N}} \neg q$ for some p and q), and, conversely, if a net is inconsistent, then there is a node p with infinite ranking.

Proposition 3.9 $\rightsquigarrow_{\text{BIN}}$ satisfies (REF), (CM), and (CT).

Proof. The proof retraces the one for proposition 3.3. \square

Proposition 3.10 $\rightsquigarrow_{\text{BIN}}$ satisfies (LE), (RW), and (Sup).

Proof. The proof retraces the one for proposition 3.5. \square

Proposition 3.11 $\rightsquigarrow_{\text{BIN}}$ satisfies (OR) and (AND).

Proof. (OR): we have $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \star t$ and $\mathcal{N} \rightsquigarrow_{\text{BIN}} q \star t$, and we have $p, q \Leftrightarrow^{\vee} s \in \mathcal{S}$. Assume $\star = \rightarrow$. The case with $\star = \not\rightarrow$ is analogous.

We have $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow t$ and $\mathcal{N} \rightsquigarrow_{\text{BIN}} q \rightarrow t$, that is, $\hat{\Delta}_{p,t} \models_{\Phi} p \supset t$ and $\hat{\Delta}_{q,t} \models_{\Phi} q \supset t$. Since $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow t$ and $\mathcal{N} \rightsquigarrow_{\text{BIN}} q \rightarrow t$, by proposition 3.2 we have at least a duct between p and t and one between q and t . So, since $p, q \Leftrightarrow^{\vee} s$, there is a duct from s to t . Moreover, since there is at least a duct between s and t and $p \Rightarrow s$ and $q \Rightarrow s$, every duct connecting s to t allows for a duct connecting p to t and q to t . Hence, $\Delta_{s,t} \subseteq \Delta_{p,t}$ and $\Delta_{s,t} \subseteq \Delta_{q,t}$ holds.

On the other hand, we have that, since $p, q \Leftrightarrow^{\vee} s$, s is connected to t by all the connections moving from p and from q . So, we have that $\Delta_{p,t} \subseteq \Delta_{s,t}$ and $\Delta_{q,t} \subseteq \Delta_{s,t}$. Hence, $\Delta_{p,t} = \Delta_{q,t} = \Delta_{s,t}$ holds and the ranking functions $\mathfrak{r}_{s,t}$, $\mathfrak{r}_{p,t}$, and $\mathfrak{r}_{q,t}$ are the same as they work with the same sets of material implications.

At the propositional level, we have that $\models_{\Phi} s \equiv (p \vee q)$ and either $\mathfrak{r}_{s,t}(p) \leq \mathfrak{r}_{s,t}(q)$ or $\mathfrak{r}_{s,t}(q) \leq \mathfrak{r}_{s,t}(p)$. Assume the former (the same reasoning applies in the other case). If $\mathfrak{r}_{s,t}(p) \leq \mathfrak{r}_{s,t}(q)$ we have that $\hat{\Delta}_{q,t} \subseteq \hat{\Delta}_{p,t}$. Moreover, the least set in the exceptionality order negating p does negate also q , and so is the least set in the exceptionality order negating t , that is, $\mathfrak{r}_{s,t}(t) = \mathfrak{r}_{s,t}(p)$, and $\hat{\Delta}_{p,t} = \hat{\Delta}_{s,t}$. So, $\hat{\Delta}_{s,t} \models_{\Phi} p \supset t$, and, since $\hat{\Delta}_{q,t} \subseteq \hat{\Delta}_{p,t}$, we have also $\hat{\Delta}_{s,t} \models_{\Phi} q \supset t$; hence by classical reasoning we obtain $\hat{\Delta}_{s,t} \models_{\Phi} (p \vee q) \supset t$, i.e., $\hat{\Delta}_{s,t} \models_{\Phi} s \supset t$. Eventually, $\mathcal{N} \rightsquigarrow_{\text{BIN}} s \rightarrow t$.

(AND): we have $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow q$ and $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow s$, and we have $s, q \Leftrightarrow^{\wedge} t \in \mathcal{S}$. Every duct connecting p to q and p to s are part of a duct connecting p to t by the addition of $s, q \Leftrightarrow^{\wedge} t$, while every duct connecting p to t is part of a duct connecting p to q (resp., p to s) by the addition of $t \Rightarrow q$ (resp., $t \Rightarrow s$). So, $\Delta_{p,q} = \Delta_{p,s} = \Delta_{p,t}$, and we obtain $\mathcal{N} \rightsquigarrow_{\text{BIN}} p \rightarrow t$. \square

Proposition 3.12 A BIN \mathcal{N} is consistent iff we do not have any node p with $\mathfrak{r}(p) = \infty$, that is, we cannot conclude both $p \rightarrow q$ and $p \not\rightarrow q$ for any pair p, q .

Proof. The proof is similar to the one for proposition 3.6. \square

Proposition 4.1. $\vdash_{\mathcal{K}}$ is a rational consequence relation containing \mathcal{K} .

Proof. $\vdash_{\mathcal{K}}$ has been obtained as the rational closure of \mathcal{K}' , so it is a rational consequence relation. Moreover, $\mathbb{R}(\mathcal{K}')$ contains \mathcal{K}' , and, since $\mathcal{K} \subseteq \mathcal{K}'$, it contains \mathcal{K} too. \square

Proposition 4.2. *Given a conditional base \mathcal{K} , $\top \vdash_{\mathcal{K}} \perp$ iff $\top \vdash \perp \in \mathbb{R}(\mathcal{K})$.*

Proof. It is sufficient to prove that, given a conditional base $\mathcal{K} = \langle \mathcal{T}, \mathcal{U} \rangle$, and obtained its extension $\mathcal{K}' = \langle \mathcal{T}, \mathcal{U}' \rangle$ (Steps 1-3), \mathcal{K} is preferentially consistent iff the conditional base \mathcal{K}' is preferentially consistent too.

From right to left the proof is immediate, since $\mathcal{K} \subseteq \mathcal{K}'$. From left to right, if \mathcal{K} is consistent that means that $\overline{\mathcal{T}} \cup \overline{\mathcal{D}} \not\vdash \perp$. Given that every conditional added in \mathcal{K}' corresponds to an implication that is classically derivable from a subset of $\overline{\mathcal{T}} \cup \overline{\mathcal{D}}$, we have that $\overline{\mathcal{T}} \cup \overline{\mathcal{D}'} \subseteq \mathbb{C}(\overline{\mathcal{T}} \cup \overline{\mathcal{D}})$ (where \mathbb{C} is the classical closure operation associated to \vDash). So we have $\mathbb{C}(\overline{\mathcal{T}} \cup \overline{\mathcal{D}'}) = \mathbb{C}(\overline{\mathcal{T}} \cup \overline{\mathcal{D}})$. Then, $\perp \notin \mathbb{C}(\Gamma_{\mathcal{K}'})$, and \mathcal{K}' is a consistent knowledge base too. \square

Proposition 5.1 $\vdash_{\mathcal{K}}$ *is a rational consequence relation containing \mathcal{K} .*

Proof. It is sufficient to show (1) that every inclusion axiom in \mathcal{T} and in \mathcal{D} is valid in $\vdash_{\mathcal{K}}$, and (2) that $\vdash_{\mathcal{K}}$ satisfies the properties characterizing a rational consequence relation.

- (1) By construction, if $C \sqsubseteq D \in \mathcal{T}$, then $\top \sqsubseteq C \supset D \in \tilde{\mathcal{T}}$ (that is, modulo logical equivalence, $\mathcal{T} \subseteq \tilde{\mathcal{T}}$). Assume $C \vdash_{\mathcal{K}} D \in \mathcal{D}$. Then we have either $\mathfrak{r}(C) = \mathfrak{r}(C \vdash_{\mathcal{K}} D) = \infty$, or $\mathfrak{r}(C) = \mathfrak{r}(C \vdash_{\mathcal{K}} D) = i$, for some $i < \infty$. In the first case we have $\top \sqsubseteq \neg C \in \tilde{\mathcal{T}}$, and so $\vDash \{C\} \cup \tilde{\Phi} \sqsubseteq D$, which implies $C \vdash_{\mathcal{K}} D$. In the second case, we have that $\vDash \delta_i \sqsubseteq C \supset D$, with δ_i being the default-assumption associated to the premise C . Hence we have $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq D$, that is, $C \vdash_{\mathcal{K}} D$.
- (2) The consequence relation $\vdash_{\mathcal{K}}$ satisfies the properties of a rational consequence relation. *(REF)* is obviously satisfied, as *(LLE)* is. For *(RW)*, assume we have $C \vdash_{\mathcal{K}} D$. That means that for the first C -consistent δ_i in Δ , $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq D$, that, given $\vDash_{\tilde{\Phi}} D \sqsubseteq E$, implies $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq E$, i.e., $C \vdash_{\mathcal{K}} E$.
 - (CT)* $C \sqcap D \vdash_{\mathcal{K}} E$ corresponds to $\vDash \{C \sqcap D\} \cup \tilde{\Phi} \cup \delta_i \sqsubseteq E$, where δ_i is the first $C \sqcap D$ -consistent formula in Δ . Analogously, $C \vdash_{\mathcal{K}} D$ means that for the first C -consistent δ_j in Δ , $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq D$. Since $\vDash C \sqcap D \sqsubseteq C$, we have that $j \leq i$, that is, $\vDash \delta_j \sqsubseteq \delta_i$. Hence, we have that $\vDash \{C \sqcap D\} \cup \tilde{\mathcal{T}} \cup \{\delta_j\} \sqsubseteq E$ and $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq D$, and, since \vDash satisfies *(CT)*, $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq E$, that is, $C \vdash_{\mathcal{K}} E$.
 - (CM)* $C \vdash_{\mathcal{K}} E$ means that for the first C -consistent δ_i in Δ , $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq E$. Analogously, $C \vdash_{\mathcal{K}} D$ means that $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq D$ for the same δ_i . Hence, δ_i is consistent with $C \sqcap D$, and, by the monotony of \vDash , we have $\vDash \{C \sqcap D\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq E$, that is, $C \sqcap D \vdash_{\mathcal{K}} E$.
 - (OR)* $C \vdash_{\mathcal{K}} E$ means that for the first C -consistent δ_i in Δ , $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq E$. Analogously, $D \vdash_{\mathcal{K}} E$ means that $\vDash \{D\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq E$ for the first D -consistent δ_j in Δ . We have three options: $j = i$, $j < i$ or $i < j$. In the first case, the default-assumption associated with $C \sqcup D$ is δ_i , and, since \vDash satisfies *OR*, $\vDash \{C \sqcup D\} \cup \tilde{\mathcal{T}} \cup \{\delta_i\} \sqsubseteq E$, that is, $C \sqcup D \vdash_{\mathcal{K}} E$. If $j < i$, we have that $\delta_j \vDash \delta_i$, and that δ_j is the first $C \sqcup D$ -consistent default in Δ . Hence we have $\vDash \{D\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq E$, and, by monotonicity, $\vDash \{C\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq E$. Since \vDash satisfies *OR*, we have $\vDash \{C \sqcup D\} \cup \tilde{\Phi} \cup \{\delta_j\} \sqsubseteq E$, that is, $C \sqcup D \vdash_{\mathcal{K}} E$.

(RM) $C \not\sim_{\mathcal{K}} \neg E$ corresponds to $\not\models \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq \neg E$, where δ_i is the first C -consistent formula in Δ . This means that E is consistent with $\{C\} \cup \tilde{\Phi} \cup \{\delta_i\}$, and so δ_i is the first $C \sqcap E$ -consistent formula in Δ . Since $\models \{C\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq D$, by the monotonicity of \models we have $\models \{C \sqcap E\} \cup \tilde{\Phi} \cup \{\delta_i\} \sqsubseteq D$, that is, $C \sqcap E \sim_{\mathcal{K}} D$.

□

Proposition 5.3 *Given a linear order of the individuals in \mathcal{K} , our procedure determines an ABox extension of \mathcal{K} . Vice-versa, every ABox extension of \mathcal{K} corresponds to the knowledge base generated by some linear order of the individuals in \mathcal{O} .*

Proof. The first statement is quite immediate. For the second statement, assume that there is a rational extension $\langle \mathcal{A}', \mathcal{T} \rangle$ of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ that cannot be generated by any sequence s of the elements of \mathcal{O} . \mathcal{A}' associates to every individual \mathbf{x} a default concept from Δ , that we indicate as $\delta^{\mathbf{x}}$.

Now, consider a generic rational extension $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$ of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ that can be generated using a sequence of elements of \mathcal{O} . The following procedure allows to define a sequence s of the elements of \mathcal{O} s.t. $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$ can be generated using s , i.e., $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle = \langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$.

Take each element of \mathcal{O} and associate to it the strongest default concept in Δ consistent with the knowledge base $\langle \mathcal{A}, \mathcal{T} \rangle$ (call it $\gamma^{\mathbf{x}}$). Look for an individual \mathbf{x} s.t. $\delta^{\mathbf{x}} = \gamma^{\mathbf{x}}$, and consider \mathbf{x} the first element of the sequence s . Update \mathcal{A} with $\mathbf{x}:\delta^{\mathbf{x}}$, and repeat the procedure, until every individual has been associated to a default formula. With this procedure we can generate a sequence over the dominion of the individuals that generates $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T} \rangle$ from $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$.

Since there is no sequence s that can generate $\langle \mathcal{A}', \mathcal{T} \rangle$, the above procedure has to fail, that is, at some point it will not be possible to associate to any \mathbf{x} a default $\gamma^{\mathbf{x}}$ s.t. $\delta^{\mathbf{x}} = \gamma^{\mathbf{x}}$. That means that, for all the remaining \mathbf{x} , $\delta^{\mathbf{x}} \neq \gamma^{\mathbf{x}}$; for each such \mathbf{x} , either $\delta^{\mathbf{x}} \models \gamma^{\mathbf{x}}$ or $\gamma^{\mathbf{x}} \models \delta^{\mathbf{x}}$. The first case is not possible, since $\langle \mathcal{A}', \mathcal{T} \rangle$ would be inconsistent ($\gamma^{\mathbf{x}}$ has to be a maximally consistent default). Hence $\gamma^{\mathbf{x}} \models \delta^{\mathbf{x}}$ and $\delta^{\mathbf{x}} \neq \gamma^{\mathbf{x}}$ for all the remaining x . In such a case, $\langle \mathcal{A}', \mathcal{T} \rangle$ would not be a rational extension of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, since we could have another consistent model with stronger defaults associated to some individuals. □

Proposition 5.4 *Given \mathcal{K} and a linear order s of the individuals in \mathcal{K} , the consequence relation \Vdash_s satisfies the properties (REF_{DL}) – (RM_{DL}).*

Proof. For *REF_{DL}*, *LLE_{DL}* and *RW_{DL}* the proof is quite immediate. For *CT_{DL}* and *CM_{DL}*, assume $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{b}:D$, that is $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{b}:D$. Hence, $\mathbf{b}:D$ is consistent with $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$; this implies that for the first individual in the sequence s , let it be \mathbf{a} , and for every $\delta_i \in \Delta$, $\mathbf{a}:\delta_i$ is consistent with $\langle \mathcal{A}, \mathcal{T} \rangle$ iff it is consistent with $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T} \rangle$, and the procedure associates to \mathbf{a} the same default formula either we start with \mathcal{A} or with $\mathcal{A} \cup \{\mathbf{b}:D\}$. The same happens for all the following individuals in the sequence s . So we have that $\langle \mathcal{A}_{\mathcal{D}}^s \cup \{\mathbf{b}:D\}, \mathcal{T} \rangle = \langle (\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s, \mathcal{T} \rangle$ and $\langle \mathcal{A}_{\mathcal{D}}^s \cup \{\mathbf{b}:D\}, \mathcal{T} \rangle \models \mathbf{a}:C$ iff $\langle (\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$. Since \models satisfies *CT* and *CM*, we have that $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$ iff $\langle (\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$, that is, $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$ iff $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$.

For *OR_{DL}*, assume that $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$, $\langle \mathcal{A} \cup \{\mathbf{b}:E\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$, and that \mathbf{b} is in the n^{th} position in the sequence s . So, for the first $n - 1$ elements of s the association with the default-formulae is the same in both the models. For \mathbf{b} , assume that the procedure

assigns $\mathbf{b}:\delta_i$ in case $\mathbf{b}:D$, and $\mathbf{b}:\delta_j$ in case $\mathbf{b}:E$. We can have either $\delta_i = \delta_j$, $\models \delta_i \sqsubseteq \delta_j$, or $\models \delta_j \sqsubseteq \delta_i$. In the first case the procedure for the assignment of the defaults continues in the same way in both the knowledge bases, and it is the same also if we have $\mathbf{b}:D \sqcup E$, that is, $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle$, $\langle \mathcal{A} \cup \{\mathbf{b}:E\}, \mathcal{T}, \mathcal{D} \rangle$, and $\langle \mathcal{A} \cup \{\mathbf{b}:D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ are completed exactly with the same defaults, obtaining, respectively, the ABoxes $(\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s = \mathcal{A}' \cup \{\mathbf{b}:D\}$, $(\mathcal{A} \cup \{\mathbf{b}:E\})_{\mathcal{D}}^s = \mathcal{A}' \cup \{\mathbf{b}:E\}$, and $(\mathcal{A} \cup \{\mathbf{b}:D \sqcup E\})_{\mathcal{D}}^s = \mathcal{A}' \cup \{\mathbf{b}:D \sqcup E\}$, for some ABox \mathcal{A}' . So we have that $\mathcal{A}' \cup \{\mathbf{b}:D\} \models \mathbf{a}:C$ and $\mathcal{A}' \cup \{\mathbf{b}:E\} \models \mathbf{a}:C$, and, since \models satisfies OR, we obtain $\mathcal{A}' \cup \{\mathbf{b}:D \sqcup E\} \models \mathbf{a}:C$, that is, $\langle (\mathcal{A} \cup \{\mathbf{b}:D \sqcup E\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$. If $\delta_i \models \delta_j$ and $\mathbf{b}:D \sqcup E$, the procedure associates to \mathbf{b} the strongest of the two defaults, that is, δ_i . Since δ_i is not consistent with $\mathbf{b}:E$, in every following consistency check the procedure will be forced to consider that $\mathbf{b}:D$ holds, and the assignment of the defaults to the individuals will proceed as in the case where $\mathbf{b}:D$ holds, and $\langle \mathcal{A} \cup \{\mathbf{b}:D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ will entail the same formulae as $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle$. Analogously, if $\delta_j \models \delta_i$, the default-assumption extension of $\langle \mathcal{A} \cup \{\mathbf{b}:D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ will correspond to the one of $\langle \mathcal{A} \cup \{\mathbf{b}:E\}, \mathcal{T}, \mathcal{D} \rangle$.

Finally, for RM_{DL} , $\mathbf{b}:D$ is consistent with $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$, so the presence of $\mathbf{b}:D$ in the knowledge base does not influence the association of the defaults to the individuals, and $\mathcal{A}_{\mathcal{D}}^s \subseteq (\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s$. Eventually, $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$ implies $\langle (\mathcal{A} \cup \{\mathbf{b}:D\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models \mathbf{a}:C$, i.e. $\langle \mathcal{A} \cup \{\mathbf{b}:D\}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$.

□

Proposition 5.7 $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \Vdash \mathbf{a}:C$ iff $\langle \mathcal{A}_{[a]}, \mathcal{T}, \mathcal{D} \rangle \Vdash_s \mathbf{a}:C$ for every ordering s of the individuals in $\mathcal{A}_{[a]}$.

Proof. The proof is quite immediate. Assume $\langle \mathcal{A}_{[a]}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash_s \mathbf{a}:C$ for some s . Let s' be a sequence of the individuals named in \mathcal{A} obtained using s as initial segment. Hence we have that $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash^{s'} \mathbf{a}:C$, the implies $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash \mathbf{a}:C$.

Now assume $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash \mathbf{a}:C$. Hence, for some sequence s $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash^s \mathbf{a}:C$. Let s' be a restriction of s to the individuals named in $\mathcal{A}_{[a]}$; then we have that $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\Vdash^{s'} \mathbf{a}:C$. □

Appendix C. Table of the Main Symbols

Since the paper considers different fields, the notation has turned out to be quite complex. We add here a table to summarize the main symbols used in the paper.

	IN/BIN	PL	DL
p, q, \dots	nodes	atoms	concept names
C, D, \dots	-	propositions	concepts
$\mathbf{a}, \mathbf{b}, \dots$	-	-	individuals
\vdash	strict conditional	strict conditional	strict conditional
\vdash	defeasible conditional	defeasible conditional	defeasible conditional
\rightsquigarrow	consequence relation	-	-
\otimes, \oplus	conjunction, disjunction	-	-
$\mathcal{C}_{p,q}^{\mathcal{N}}$	links in the courses/ducts between p and q	links in the ducts between p and q	links in the ducts between p and q
$\Delta_{p,q}^{\mathcal{N}}$	materialisations of the links in $\mathcal{C}_{p,q}^{\mathcal{N}}$	materialisations of the links in $\mathcal{C}_{p,q}^{\mathcal{N}}$	materialisations of the links in $\mathcal{C}_{p,q}^{\mathcal{N}}$
$\hat{\Delta}_{p,q}^{\mathcal{N}}$	conditionals in $\Delta_{p,q}^{\mathcal{N}}$ as exceptional as p	conditionals in $\Delta_{p,q}^{\mathcal{N}}$ as exceptional as p	conditionals in $\Delta_{p,q}^{\mathcal{N}}$ as exceptional as p

References

- Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, pp. 510–530.
- Baader, F., Calvanese, D., McGuinness, D., Nardi, D., & Patel-Schneider, P. F. (Eds.). (2003). *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.
- Baader, F., & Hollunder, B. (1993). How to prefer more specific defaults in terminological default logic. In *Proceedings of IJCAI-93*, pp. 669–674. Morgan Kaufmann Publishers.
- Bochman, A. (2001). *A logical theory of nonmonotonic inference and belief change*. Springer-Verlag.
- Bonatti, P. A., Lutz, C., & Wolter, F. (2009). The complexity of circumscription in description logic. *Journal of Artificial Intelligence Research*, 35, pp. 717–773.
- Bonatti, P. A., Faella, M., & Sauro, L. (2010). \mathcal{EL} with default attributes and overriding. In Patel-Schneider, P. F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J. Z., Horrocks, I., & Glimm, B. (Eds.), *International Semantic Web Conference (1)*, Vol. 6496 of *Lecture Notes in Computer Science*, pp. 64–79. Springer.
- Bonatti, P. A., Faella, M., & Sauro, L. (2011a). Adding default attributes to \mathcal{EL}^{++} . In Burgard, W., & Roth, D. (Eds.), *Proceedings of AAAI-11*. AAAI Press.
- Bonatti, P. A., Faella, M., & Sauro, L. (2011b). Defeasible inclusions in low-complexity DLs. *Journal of Artificial Intelligence Research*, 42, pp. 719–764.
- Bonatti, P. A., Faella, M., & Sauro, L. (2011c). On the complexity of \mathcal{EL} with defeasible inclusions. In *Proceedings of IJCAI-11*, pp. 762–767. AAAI Press/IJCAI.
- Brewka, G., & Augustin, D. S. (1987). The logic of inheritance in frame systems. In *Proceedings of IJCAI-87*, pp. 483–488. Morgan Kaufmann Publishers.
- Britz, K., Casini, G., Meyer, T., Moodley, K., & Varzinczak, I. (2013). Ordered Interpretations and Entailment for Defeasible Description Logics. Tech. rep., CAIR, CSIR Meraka and UKZN, South Africa.
- Britz, K., Heidema, J., & Meyer, T. A. (2008). Semantic preferential subsumption. In Brewka, G., & Lang, J. (Eds.), *Proceedings of KR-08*, pp. 476–484. AAAI Press.
- Britz, K., Meyer, T., & Varzinczak, I. J. (2011). Semantic foundation for preferential description logics. In Wang, D., & Reynolds, M. (Eds.), *Australasian Conference on Artificial Intelligence*, Vol. 7106 of *Lecture Notes in Computer Science*, pp. 491–500. Springer.
- Casini, G., Meyer, T., Moodley, K., & Varzinczak, I. (2013). Nonmonotonic reasoning in description logics. Rational closure for the ABox. In *Proceedings of DL-13*, pp. 77–90. CEUR Workshop Proceedings.
- Casini, G., & Straccia, U. (2010). Rational closure for defeasible description logics. In Janhunen, T., & Niemelä, I. (Eds.), *Proceedings of JELIA-10*, Vol. 6341 of *Lecture Notes in Computer Science*, pp. 77–90. Springer.

- Casini, G., & Straccia, U. (2011). Defeasible inheritance-based description logics. In *Proceedings of IJCAI-11*, pp. 813–818.
- Cormen, T. H., Stein, C., Rivest, R. L., & Leiserson, C. E. (2001). *Introduction to Algorithms* (2nd edition). McGraw-Hill Higher Education.
- Donini, F. M., & Massacci, F. (2000). EXP-TIME tableaux for \mathcal{ALC} . *Artificial Intelligence*, 124(1), pp. 87–138.
- Donini, F. M., Nardi, D., & Rosati, R. (2002). Description logics of minimal knowledge and negation as failure. *Transactions on Computational Logic*, 3(2), pp. 177–225.
- Freund, M. (1998). Preferential reasoning in the perspective of Poole default logic. *Artificial Intelligence*, 98(1-2), pp. 209–235.
- Gabbay, D. M., & Schlechta, K. (2009). Defeasible inheritance systems and reactive diagrams. *Logic Journal of the IGPL*, 17(1), pp. 1–54.
- Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. L. (2012a). A minimal model semantics for nonmonotonic reasoning. In *Proceedings of JELIA-12*, Vol. 7519 of *Lecture Notes in Computer Science*, pp. 228–241. Springer.
- Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. L. (2012b). A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, 195, pp. 165–202.
- Giordano, L., Olivetti, N., Gliozzi, V., & Pozzato, G. L. (2009). $\mathcal{ALC}+T$: a preferential extension of description logics. *Fundam. Inform.*, 96(3), 341–372.
- Grimm, S., & Hitzler, P. (2009). A preferential tableaux calculus for circumscriptive \mathcal{ALCO} . In *Proceedings of RR-09*, pp. 40–54. Springer-Verlag.
- Horty, J. F. (1994). Some direct theories of nonmonotonic inheritance. In *Handbook of logic in artificial intelligence and logic programming: nonmonotonic reasoning and uncertain reasoning*, Vol. 3, pp. 111–187. Oxford University Press.
- Horty, J. F., & Thomason, R. H. (1990). Boolean extensions of inheritance networks. In *Proceedings of AAAI-90*, pp. 633–639. AAAI Press.
- Horty, J. F., Thomason, R. H., & Touretzky, D. S. (1987). A skeptical theory of inheritance in nonmonotonic semantic networks. In *Proceedings of AAAI-87*. AAAI Press.
- Knorr, M., Alferes, J. J., & Hitzler, P. (2011). Local closed world reasoning with description logics under the well-founded semantics. *Artificial Intelligence*, 175(9-10), 1528–1554.
- Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2), pp. 167–207.
- Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail?. *Artificial Intelligence*, 55(1), pp. 1–60.
- Lukasiewicz, T. (2008). Expressive probabilistic description logics. *Artificial Intelligence*, 172(6-7), pp. 852–883.
- Makinson, D. (1994). General patterns in nonmonotonic reasoning. In *Handbook of logic in artificial intelligence and logic programming: nonmonotonic reasoning and uncertain reasoning*, Vol. 3, pp. 35–110. Oxford University Press.

- Makinson, D. (2005). *Bridges from Classical to Nonmonotonic Logic*. King's College Publications.
- Makinson, D., & Schlechta, K. (1991). Floating conclusions and zombie paths. *Artificial Intelligence*, 48, pp. 199–209.
- Poole, D. (1988). A logical framework for default reasoning. *Artificial Intelligence*, 36(1), 27–47.
- Quantz, J., & Royer, V. (1992). A preference semantics for defaults in terminological logics. In *Proceedings of KR-92*, pp. 294–305.
- Rott, H. (2001). *Change, Choice and Inference: a study of belief revision and nonmonotonic reasoning*. Oxford University Press.
- Sandewall, E. (1986). Nonmonotonic inference rules for multiple inheritance with exceptions. In *Proceedings of the IEEE-86*, pp. 1345–1353.
- Sandewall, E. (2010). Defeasible inheritance with doubt index and its axiomatic characterization. *Artificial Intelligence*, 18(174), pp. 1431–1459.
- Slechta, K. (2004). *Coherent Systems*. Elsevier.
- Shoham, Y. (1988). *Reasoning about change: time and causation from the standpoint of artificial intelligence*. MIT Press.
- Simonet, G. (1996). On sandewall's paper: Nonmonotonic inference rules for multiple inheritance with exceptions. *Artificial Intelligence*, 86, pp. 359–374.
- Straccia, U. (1993). Default inheritance reasoning in hybrid KL-ONE style logics. *Proceedings of IJCAI-93*, 676–681.
- Thomason, R. H. (1992). NETL and subsequent path-based inheritance theories. In Lehmann, F. (Ed.), *Semantic Networks in Artificial Intelligence*, pp. 179–204. Pergamon Press.
- Touretzky, D. S. (1986). *The mathematics of inheritance systems*. Pitman.
- Touretzky, D. S., Horty, J. F., & Thomason, R. H. (1987). A clash of intuitions: the current state of nonmonotonic multiple inheritance systems. In *Proceedings of IJCAI-87 - Vol. 1*, pp. 476–482. Morgan Kaufmann Publishers.
- Touretzky, D. S., Thomason, R. H., & Horty, J. F. (1991). A skeptic's menagerie: conflictors, preemptors, reinstates, and zombies in nonmonotonic inheritance. In *Proceedings of IJCAI-91*, pp. 478–483. Morgan Kaufmann Publishers.