

On the Logical Properties of the Description Logic \mathcal{DL}^N (Extended abstract)*

Piero A. Bonatti and Luigi Sauro
 Università degli Studi di Napoli Federico II
 {piero.bonatti, luigi.sauro}@unina.it

Abstract

\mathcal{DL}^N is a recent nonmonotonic description logic, designed for satisfying independently proposed knowledge engineering requirements, and for removing some recurrent drawbacks of traditional nonmonotonic semantics. In this paper we study the logical properties of \mathcal{DL}^N and their relationships with the KLM postulates. We use various versions of the KLM postulates to deepen the comparison with related work, and illustrate the different tradeoffs between opposite expressivity requirements adopted by each approach.

1 Introduction

Recently, in [Bonatti *et al.*, 2015a], a new nonmonotonic description logic called \mathcal{DL}^N has been introduced with the goal of supporting ontology authoring by means of nonmonotonic reasoning. \mathcal{DL}^N aims at removing some recurrent drawbacks of traditional nonmonotonic DLs¹ such as: (i) *inheritance blocking*, a drawback of preferential semantics and rational closure; (ii) undesired *closed world assumption effects*, that affect circumscription, typicality logic and some probabilistic logics; (iii) the inability to specify whether roles shall range over normal/prototypical individuals or not, that affects most nonmonotonic DLs. Moreover, \mathcal{DL}^N adopts a novel conflict resolution mechanism that helps in detecting unresolved conflicts between mutually inconsistent defaults. Since unresolved conflicts frequently correspond to missing knowledge, highlighting such conflicts constitutes an important support to knowledge base debugging and validation. A further useful property of \mathcal{DL}^N is that it can be translated into classical DLs, so that its implementations can rely on mature and well-optimized inference engines. The translation can be computed in polynomial time and does not involve complex constructs, so \mathcal{DL}^N preserves the tractability of the two major low-complexity families of DLs (\mathcal{EL} and \mathcal{DL} -lite, that correspond to the EL and QL profiles of OWL2, respectively). The experiments in [Bonatti *et al.*, 2015a; 2015b] show unparalleled scalability properties over large

knowledge bases, with more than 10^5 axioms. The relationships between the above features and the knowledge engineering requirements that have been independently introduced in the fields of biomedical ontologies and declarative policy languages have been extensively discussed in [Bonatti *et al.*, 2015a].

The main goal of this paper is studying the logical properties of \mathcal{DL}^N 's consequence relation, defeasible axioms and normal instances. Our analysis includes a comparison of \mathcal{DL}^N 's inferences with verbatim and internalized versions of the KLM postulates [Kraus *et al.*, 1990]. For our purposes, these postulates are not necessarily desiderata, due to the loose correspondence between their motivations and \mathcal{DL}^N 's goals and semantics (cf. [Bonatti *et al.*, 2015a; Bonatti and Sauro, 2017]). However, we regard them as a useful technical tool for profiling and comparing the behavior of different logics, since the validity (or non-validity) of the postulates has been extensively investigated in most nonmonotonic logics.

In [Bonatti and Sauro, 2017] we fix also some technical problems affecting the original definition of \mathcal{DL}^N ; here we skip these technical details due to space limitations.

2 The Basics of \mathcal{DL}^N

Let \mathcal{DL} be any classical description logic language and let \mathcal{DL}^N be the extension of \mathcal{DL} with a new concept name NC for each \mathcal{DL} concept C . The new concepts are called *normality concepts* and denote the standard instances of C . Besides the usual axioms of \mathcal{DL} , called *classical* or *strong*, \mathcal{DL}^N further supports new axioms called *defeasible inclusions* (DIs, for short) that are expressions $C \sqsubseteq_n D$ where C is a \mathcal{DL} concept and D a \mathcal{DL}^N concept. The informal meaning of $C \sqsubseteq_n D$, roughly speaking, is: “by default, the standard instances of any concept E should satisfy $\neg C \sqcup D$, unless stated otherwise”, that is, unless some higher priority (possibly strong) axioms entail that the standard instances of E belong to $C \sqcap \neg D$; in that case, we say that $C \sqsubseteq_n D$ is *overridden in E* . The standard instances of a concept C (i.e., the members of NC) are required to satisfy all the DIs that are not overridden in C . We will call the set of DIs satisfied in NC the *prototype* associated to C .

DIs are prioritized by a strict partial order \prec . Such priority order is part of the knowledge base. It may be based

*This is an extended abstract of [Bonatti and Sauro, 2017].

¹Abbreviation for “description logics”.

$Mamm \sqsubseteq_n \exists lungs$	$Mamm \sqsubseteq_n \neg \exists fins$
$SeaAnim \sqsubseteq_n \exists fins$	$Whales \sqsubseteq Mamm \sqcap SeaAnim$
$Dolphins \sqsubseteq Mamm \sqcap SeaAnim \sqcap \exists fins$	

 Table 1: The \mathcal{KB} formalizing Example 1

on *specificity*, to make the defeasible properties of a class C override those of its superclasses, but it may also be based on time, to make more recent axioms override older axioms (as it is customary in law). If $\delta_1 \prec \delta_2$, then δ_1 has higher priority than δ_2 . \mathcal{DL}^N solves automatically only the conflicts that can be settled using \prec . Any other conflict shall be resolved by the knowledge engineer (typically by adding more DIs). Unresolved conflicts cause the affected concept C to have no standard instances, that is, $NC \sqsubseteq \perp$ is entailed by the knowledge base (in symbols: $\mathcal{KB} \approx NC \sqsubseteq \perp$). Thus unresolved conflicts can be detected by simple concept satisfiability tests, while all the other logics just hide the conflicts.

The papers [Bonatti *et al.*, 2015a; 2015b; Bonatti and Sauro, 2017] contain examples inspired by real applications, but here – due to space limitations – we resort to an artificial example that illustrates several features of \mathcal{DL}^N at once.

Example 1 Let us formalize a domain where “by default, mammals have lungs and not fins, by default sea animals have fins, whales are both mammals and sea animals, dolphins are both mammals and sea animals and have fins”. The axiomatization in \mathcal{DL}^N is reported in Table 1. The priority relation on DIs is specificity.

Both *Dolphins* and *Whales* are in the intersection of *Mamm* and *SeaAnim*, that have conflicting default properties concerning fins. Since *Mamm* and *SeaAnim* are incomparable, specificity does not settle the conflict.

However, dolphins are explicitly asserted to have fins; this assertion overrides the second default property of mammals thereby resolving the conflict. Still, Dolphins inherit the other default property of mammals (since \mathcal{DL}^N is not affected by *inheritance blocking*) so the attributes of standard dolphins are captured by the following inference, stating in formal terms that *standard dolphins have both fins and lungs*:

$$\mathcal{KB} \approx N \text{ Dolphins} \sqsubseteq \exists fins \sqcap \exists lungs.$$

The definition of *Whales*, instead, is incomplete: the knowledge engineer forgot to settle the conflict. However, the need for additional knowledge can be easily identified by looking for inconsistent normality concepts. In \mathcal{DL}^N (where unresolved conflicts are not silently removed because a DI can be overridden only by axioms with *strictly higher* priority) whales inherit *all* the (mutually inconsistent) default properties of mammals and sea animals, therefore:

$$\mathcal{KB} \approx N \text{ Whales} \sqsubseteq \perp.$$

In other words, the standard concept consistency tests carried out during KB validation point out also unresolved conflicts.

In some logics, like circumscribed DLs and typicality DLs, the fact that dolphins are exceptional mammals would cause

Name	Rule schema	Valid in \mathcal{DL}^N
REF	$\frac{\alpha \in \mathcal{KB}}{\mathcal{KB} \approx \alpha}$	✓
CT	$\frac{\mathcal{KB} \approx \alpha \quad \mathcal{KB} \cup \{\alpha\} \approx \gamma}{\mathcal{KB} \approx \gamma}$	✓
CM	$\frac{\mathcal{KB} \approx \alpha \quad \mathcal{KB} \approx \gamma}{\mathcal{KB} \cup \{\alpha\} \approx \gamma}$	✓
LLE	$\frac{\mathcal{KB} \cup \{\alpha\} \approx \gamma \quad \models \alpha \equiv \beta}{\mathcal{KB} \cup \{\beta\} \approx \gamma}$	✓
RW	$\frac{\mathcal{KB} \approx \alpha \quad \alpha \models \gamma}{\mathcal{KB} \approx \gamma}$	✓
OR	$\frac{\mathcal{KB} \cup \{\alpha\} \approx \gamma \quad \mathcal{KB} \cup \{\beta\} \approx \gamma}{\mathcal{KB} \cup \{\alpha \vee \beta\} \approx \gamma}$	
RM	$\frac{\mathcal{KB} \approx \gamma \quad \mathcal{KB} \not\approx \neg \alpha}{\mathcal{KB} \cup \{\alpha\} \approx \gamma}$	

Nonstandard DL axioms $\alpha \vee \beta$, $\neg \beta$ can be simulated with the universal role;
 \approx denotes the nonmonotonic consequence relation of \mathcal{DL}^N ;
 \models denotes classical inference.
All these rules hold when there are no unresolved conflicts.

 Table 2: The KLM postulates in \mathcal{DL}^N

the set of such exceptional individuals to be minimized. So, in the direct equivalents of the above knowledge base, *Dolphins* would be inconsistent ($Dolphins \sqsubseteq \perp$). On the contrary, in \mathcal{DL}^N neither *Dolphins* nor $N \text{ Dolphins}$ are inconsistent. If we further asserted that Moby is a whale, then circumscribed DLs and typicality DLs would infer $Whales \equiv \{Moby\}$ while \mathcal{DL}^N would not. ■

3 Relationships with the KLM Postulates

In their seminal papers, Makinson [Makinson, 1988] and Kraus, Lehmann, and Magidor [Kraus *et al.*, 1990; Lehmann and Magidor, 1992; Lehmann, 1995] argued that in order to reason about what normally holds in the world, it is desirable to make nonmonotonic consequence relations closed under certain properties, called *KLM postulates*, from the initials of their authors. Even if \mathcal{DL}^N has not been built with the KLM postulates in mind (because it has different goals), it often behaves in a quite similar way. In \mathcal{DL}^N there exist several analogues of the postulates; the verbatim instantiation of the original, meta-level postulates – in \mathcal{DL}^N ’s terminology – is illustrated in Table 2. With respect to this version of the postulates, \mathcal{DL}^N is *cumulative*, i.e. it satisfies the first five rules. The cumulativity property shows that DIs significantly differ from default rules, since default logic is not cumulative.

If each consistent \mathcal{DL} concept has at least one standard instance (i.e. all unresolved conflicts have been removed), then \mathcal{DL}^N satisfies all postulates – that is, its consequence relation is *rational*. The same precondition is hardwired in the semantics of typicality logics through the smoothness condition, and a similar assumption was made in a first-order extension of the postulates [Lehmann and Magidor, 1990]. This is

Name	Rule schema	Valid in \mathcal{DL}^N
REF _n	$\frac{}{C \sqsubseteq_n C}$	✓
CT _n	$\frac{C \sqsubseteq_n D \quad C \sqcap D \sqsubseteq_n E}{C \sqsubseteq_n E}$	
CM _n	$\frac{C \sqsubseteq_n D \quad C \sqsubseteq_n E}{C \sqcap D \sqsubseteq_n E}$	
LLE _n	$\frac{C \sqsubseteq_n E \quad S \models C \equiv D}{D \sqsubseteq_n E}$	partly
RW _n	$\frac{C \sqsubseteq_n D \quad S \models D \sqsubseteq E}{C \sqsubseteq_n E}$	
OR _n	$\frac{C \sqsubseteq_n E \quad D \sqsubseteq_n E}{C \sqcup D \sqsubseteq_n E}$	partly
RM _n	$\frac{C \sqsubseteq_n E \quad C \not\sqsubseteq_n \neg D}{C \sqcap D \sqsubseteq_n E}$	

S is the strong part of the knowledge base.
 All these rules hold when their premises are not overridden.

Table 3: Internalized KLM postulates for DIs

incompatible with \mathcal{DL}^N 's new conflict highlighting behavior. \mathcal{DL}^N does not only facilitate the identification of unresolved conflicts; it becomes also rational after such conflicts have been removed, thus one can satisfy the KLM axioms without giving up the knowledge base debugging aids of \mathcal{DL}^N .

Since OR and RM may be violated when a concept has an inconsistent prototype, an interesting open question is *whether some KLM postulates are incompatible with the novel conflict highlighting mechanism of \mathcal{DL}^N* .

In the full paper, we prove also that the rational closure of DLs itself does *not* satisfy the equivalent of OR in Table 2, but only an internalized version thereof, illustrated below. This raises an interesting research question: *to what extent do the logics based on internalized KLM postulates satisfy their original, meta-level version?*

Several logics internalize the nonmonotonic consequence relation and push the KLM postulates to the object level (e.g. [Casini and Straccia, 2010; 2011; 2013; Giordano *et al.*, 2013; 2015]). A similar internalized version is reported in Table 3. This version of the KLM axioms is clearly affected by overriding. For instance, if the second premise of CM_n were overridden, then there would be no logical ground for supporting the conclusion. However, if the premises are not overridden, then *all the postulates of Table 3 are valid in \mathcal{DL}^N* . It is very interesting to note that a similar phenomenon can be observed in Lehmann's account of default reasoning [Lehmann, 1995]. There, the postulates hold only if overridden defaults are ignored. The second interesting remark is that two of these postulates unconditionally hold in most practically interesting cases: (i) the OR rule holds if the priority relation is specificity; (ii) LLE (left logical equivalence) holds whenever the priority relation treats logically equivalent DIs in the same way (like specificity does).

Table 4 reports another internalized version of the postulates, analogous to those satisfied by typicality logics [Gior-

Name	Rule schema	Valid in \mathcal{DL}^N
REF ^N	$\frac{}{NC \sqsubseteq C}$	✓
CT ^N	$\frac{NC \sqsubseteq D \quad N(C \sqcap D) \sqsubseteq E}{NC \sqsubseteq E}$	partly
CM ^N	$\frac{NC \sqsubseteq D \quad NC \sqsubseteq E}{N(C \sqcap D) \sqsubseteq E}$	partly
LLE ^N	$\frac{NC \sqsubseteq E \quad C \equiv D}{ND \sqsubseteq E}$	
RW ^N	$\frac{NC \sqsubseteq D \quad D \sqsubseteq E}{NC \sqsubseteq E}$	✓
OR ^N	$\frac{NC \sqsubseteq E \quad ND \sqsubseteq E}{N(C \sqcup D) \sqsubseteq E}$	partly
RM ^N	$\frac{NC \sqsubseteq E \quad NC \not\sqsubseteq \neg D}{N(C \sqcap D) \sqsubseteq E}$	partly

C, *D*, and *E* range over \mathcal{DL} concepts.
 All the above rules hold if N occurs only in the queries and there are no unresolved conflicts.

Table 4: Candidate inference rules for N inspired by KLM postulates

dano *et al.*, 2013; 2015]. In the full paper, we show that typicality DLs satisfy these postulates only because the normality criterion is assumed to be *concept-independent* (i.e. if John is more typical than Mary as a parent, then he must also be more typical than Mary as a driver, as a worker, as a tax payer, and so on). \mathcal{DL}^N does not embrace this strong assumption and – for this reason – it does not universally satisfy CT^N and CM^N.

So an interesting open question is *whether the postulates of Table 4 can possibly be satisfied by a logic that does not rely on a unique, concept independent normality relation*.

CM^N interferes also with inconsistent prototypes (and the characteristic way \mathcal{DL}^N deals with unresolved conflicts):

Example 2 Consider the simple N-free \mathcal{KB} consisting of:

$$D \sqsubseteq \neg E \quad (1)$$

$$C \sqsubseteq_n D \quad (2)$$

$$C \sqsubseteq_n E. \quad (3)$$

Here the defeasible inclusions (2) and (3) together contradict the strong inclusion (1), but they have the same priority, therefore the conflict cannot be resolved in *C*, and its prototype is inconsistent (i.e. $NC \sqsubseteq \perp$). On the contrary, the instances of $C \sqcap D$, by definition, satisfy *D* (be they normal or not); then they belong to $\neg E$, by (1), so (3) is overridden and $N(C \sqcap D)$ is consistent. This shows that CM^N does not hold. If it were applied, then $N(C \sqcap D)$ would be made inconsistent, too, which is difficult to justify: it is not clear why the strong facts $N(C \sqcap D) \sqsubseteq D$ and (1) should not override (3) in the prototype of $C \sqcap D$. ■

If N does not explicitly occur in \mathcal{KB} (which in practice means, as explained in the full paper, that *role fillers are not required to be normal*, similarly to what inevitably happens in rational closure and default DLs) then all postulates in Table 4 hold

Name	Axiom schema	Sound in \mathcal{DL}^N
neg 1	$N \neg C \sqsubseteq \neg NC$	✓
neg 2	$N \neg C \sqsupseteq \neg NC$	undesirable
and 1	$NC \sqcap ND \sqsubseteq N(C \sqcap D)$	questionable
and 2	$NC \sqcap ND \sqsupseteq N(C \sqcap D)$	undesirable
or 1	$NC \sqcup ND \sqsubseteq N(C \sqcup D)$	undesirable
or 2	$NC \sqcup ND \sqsupseteq N(C \sqcup D)$	questionable

C and *D* range over \mathcal{DL} concepts

 Table 5: Candidate axioms relating N with boolean operators

but CM^N , that holds only if unresolved conflicts are removed first (cf. Example 2). So a fair comparison of \mathcal{DL}^N with the rational closure of DLs shows that the two logics are very similar, under the expressiveness restrictions of the latter.

Given these results, it is not surprising that \mathcal{DL}^N returns the standard expected conclusions in most of the artificial examples occurring in the literature. The only differences are due to the unique unresolved-conflict handling strategy of \mathcal{DL}^N . See the full paper for extensive, detailed examples on a set of benchmarks provided by Sandewall.

In the full paper we analyze also the interplay of N with the standard boolean operators. Table 5 illustrates some natural candidates. Not all of them make sense. Normality concepts satisfy the following natural axiom schema:

$$NC \sqsubseteq C \quad (4)$$

which is *strong*, that is, it cannot be overridden. Axiom **neg 1** follows from (4) and its contrapositive. It is easy to see that **neg 2** is undesirable, instead. Together with (4) and **neg 1** it implies $NC \equiv C$, that is, there could be no exceptional individuals, and there would be no difference between defeasible inclusions and strong axioms.

The axioms for \sqcap and \sqcup violate the goal of representing exceptions, too. Many motivational examples in the literature contradict **and 2** and **or 1** (see the full paper for examples). Accordingly, neither **and 2** nor **or 1** are valid in \mathcal{DL}^N .

Axioms **and 1** and **or 2** are more controversial. They are valid in some logics, but we believe that a flexible nonmonotonic logic should not always satisfy them, for the reasons explained below.

Example 3 (Drawbacks of and 1) Suppose the following facts hold. Most students (denoted by concept S) are not older than 25. Most employees (denoted by E) have a family income that exceeds €1000 per month. Most working students ($S \sqcap E$) are older than 25, being slowed down by their job, and their family income is less than €1000 per month. It is natural to assert that, by default, the instances of NS are not older than 25, that the family income of the instances of NE is at least €1000, and that the instances of $N(S \sqcap E)$ are older than 25 and their family income is less than €1000. Note that

the default properties of $N(S \sqcap E)$, in this case, are incompatible with those of NS and NE , so this natural approach at encoding the example is not compatible with **and 1**, that would force the instances of $N(S \sqcap E)$ to contain the members of $NS \sqcap NE$, that are younger than 25 and have a family income greater than €1000. ■

The example showing the drawbacks of **and 2** is not reported here due to space limitations. In typicality DLs, **and 1** and **or 2** are strong axioms, enforced by the underlying monotonic logic $\mathcal{ALC} + \mathbf{T}$. In the full paper (cf. Appendix B) we show that they are valid because the models of typicality logics adopt a single, concept-independent normality relation (as discussed above).

On the contrary, in \mathcal{DL}^N , **and 1** and **or 2** are not valid. Should they be appropriate in some application, they could be enforced by asserting them explicitly in the knowledge base. This is an interesting method for specializing \mathcal{DL}^N and tuning the set of its valid axioms when needed, so as to obtain maximum flexibility.

4 Conclusions

\mathcal{DL}^N is a flexible logic, where axioms like **and 1** and **or 2** – that seem not to be universally desirable – are not forced by the semantics and can be asserted if and when needed. The possibility of choosing different priority relations on DIs provides further flexibility and expressiveness. Even if the semantics of \mathcal{DL}^N has nothing to do with the typical foundations of the KLM postulates (like ranked interpretation and the exceptionality rankings of rational closure), still \mathcal{DL}^N satisfies the postulates to a significant extent. After all unresolved conflicts have been removed, all the meta-level postulates are satisfied. The internalized version for DIs is satisfied whenever the premises are not overridden, as in Lehmann’s default logic (that introduced the lexicographic closure). The internalized postulates for N (Table 4) are satisfied only under stronger conditions, but the logics that currently satisfy those postulates adopt the debatable assumption that there is a unique, concept-independent notion of normality, that forces also the questionable axioms in Table 5. For all the above reasons, \mathcal{DL}^N constitutes a particularly nice tradeoff between the purely logical KLM desiderata, and the knowledge engineering needs analyzed in [Bonatti *et al.*, 2015a].

The analysis of \mathcal{DL}^N in terms of the KLM postulates provides also a fine-grained comparison with other nonmonotonic DLs (see also Sec. 7 of the full paper).

The analysis of why some KLM postulates are not unconditionally satisfied in \mathcal{DL}^N led to the formulation of several research questions, some of which are currently being investigated by other authors. For example, the unique, concept-independent normality relation of ranked interpretations is now being replaced with multiple normality relations in ongoing works such as [Giordano and Gliozzi, 2018; Britz and Varzinczak, 2017]. Hopefully, these investigations will tell whether the KLM postulates can be enforced in a semantics that supports multiple, contextualized normality criteria, thereby addressing Example 3.

References

- [Bonatti and Sauro, 2017] Piero A. Bonatti and Luigi Sauro. On the logical properties of the nonmonotonic description logic \mathcal{DL}^N . *Artif. Intell.*, 248:85–111, 2017.
- [Bonatti *et al.*, 2015a] Piero A. Bonatti, Marco Faella, Iliana Petrova, and Luigi Sauro. A new semantics for overriding in description logics. *Artif. Intell.*, 222:1–48, 2015.
- [Bonatti *et al.*, 2015b] Piero A. Bonatti, Iliana Mineva Petrova, and Luigi Sauro. Optimizing the computation of overriding. In *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part I* [2015], pages 356–372.
- [Britz and Varzinczak, 2017] Katarina Britz and Ivan José Varzinczak. Towards defeasible SROIQ. In Alessandro Artale, Birte Glimm, and Roman Kontchakov, editors, *Proceedings of the 30th International Workshop on Description Logics, Montpellier, France, July 18-21, 2017.*, volume 1879 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2017.
- [Casini and Straccia, 2010] Giovanni Casini and Umberto Straccia. Rational closure for defeasible description logics. In Tomi Janhunen and Ilkka Niemelä, editors, *JELIA*, volume 6341 of *Lecture Notes in Computer Science*, pages 77–90. Springer, 2010.
- [Casini and Straccia, 2011] Giovanni Casini and Umberto Straccia. Defeasible inheritance-based description logics. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 813–818. IJCAI/AAAI, 2011.
- [Casini and Straccia, 2013] Giovanni Casini and Umberto Straccia. Defeasible inheritance-based description logics. *J. Artif. Intell. Res. (JAIR)*, 48:415–473, 2013.
- [DBL, 2015] *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part I*, volume 9366 of *Lecture Notes in Computer Science*. Springer, 2015.
- [Giordano and Gliozzi, 2018] Laura Giordano and Valentina Gliozzi. Reasoning about multiple aspects in dls: Semantics and closure construction. *CoRR*, abs/1801.07161, 2018.
- [Giordano *et al.*, 2013] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. A non-monotonic description logic for reasoning about typicality. *Artif. Intell.*, 195:165–202, 2013.
- [Giordano *et al.*, 2015] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artif. Intell.*, 226:1–33, 2015.
- [Kraus *et al.*, 1990] Sarit Kraus, Daniel J. Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 44(1-2):167–207, 1990.
- [Lehmann and Magidor, 1990] Daniel J. Lehmann and Menachem Magidor. Preferential logics: the predicate calculus case. In Rohit Parikh, editor, *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, March 1990*, pages 57–72. Morgan Kaufmann, 1990.
- [Lehmann and Magidor, 1992] Daniel J. Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.
- [Lehmann, 1995] Daniel J. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
- [Makinson, 1988] David Makinson. General theory of cumulative inference. In Michael Reinfrank, Johan de Kleer, Matthew L. Ginsberg, and Erik Sandewall, editors, *Non-Monotonic Reasoning, 2nd International Workshop, Grasse, FRG, June 13-15, 1988, Proceedings*, volume 346 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 1988.