

A New Semantics for Overriding in Description Logics (Extended abstract)*

Piero A. Bonatti, Marco Faella, Iliana M. Petrova, Luigi Sauro

Università degli studi di Napoli Federico II

piero.bonatti@unina.it

Abstract

Nonmonotonic inferences are not yet supported by Description Logic technology, although their potential usefulness is widely recognized. Lack of support to nonmonotonic reasoning is due to a number of issues related to expressiveness, computational complexity, and optimizations. This work contributes to the practical support of nonmonotonic reasoning in description logics by introducing a new semantics designed to address knowledge engineering needs. The formalism is validated through extensive comparison with the other nonmonotonic DLs, and systematic scalability tests.

1 Introduction

Many modern applications of description logics (DLs, for short), such as biomedical ontologies and semantic web policies, provide fresh motivations for extending DLs with nonmonotonic inferences. Some recent examples stemming from the biomedical domain are discussed in [Rector, 2004; Stevens *et al.*, 2007]. There, the goal of supporting default attributes and exceptions is deemed important enough to look for alternative representation methods, based on classical reasoning and ontology design patterns. However, these solutions do not scale to more complex examples with multiple exception dimensions, as discussed in [Rector, 2004]: The number of additional concepts introduced by the patterns may grow exponentially. Moreover, such auxiliary concepts are defined using computationally expensive constructs such as disjunction. So, even if the given knowledge base belongs to some low-complexity fragment (such as the OWL2 profiles), its nonmonotonic extension is generally not tractable.

Nonmonotonic DLs natively support default inferences and exceptions. However none of the standard nonmonotonic semantics produces exactly the set of expected consequences, and this can be verified on a range of rather simple examples. Some of the major known drawbacks are the following:

Inheritance blocking: In several nonmonotonic logics a concept with exceptional properties inherits *none* of the default properties of its superclasses. For instance penguins, that are exceptional birds because they do not fly, do not inherit any of the other default properties of birds, such as having wings and feathers.

Undesired CWA effects: In some nonmonotonic DLs, an exceptional concept is shrunk to the individuals that explicitly belong to it; it may even become inconsistent. Using the penguin example again, by default no penguins exist unless explicitly stated otherwise. After asserting that the individual Opus is a penguin, the concept *Penguin* becomes the singleton {Opus}.

Limited control on role ranges: In most nonmonotonic DLs, the knowledge engineer cannot specify whether a role should range only over normal individuals or not. In the DLs based on default logic, rational closure, and variants thereof, default properties never apply to role fillers.

Silent removal of unresolved conflicts: Very frequently, unresolved conflicts between nonmonotonic assertions are a symptom of a gap in the axiomatization. The correct resolution of such conflicts is typically domain dependent and should require human intervention (see [Bonatti *et al.*, 2015a; 2015b] for a detailed discussion of this issue). Most nonmonotonic logics hide such conflicts (i.e. they do not have any visible consequence) thereby hindering their identification – a necessary step for the validation and correction of knowledge bases.

Moreover, the computational complexity of nonmonotonic DLs is almost always higher than the complexity of the corresponding classical DL, and the tractability of the OWL 2 profiles is not preserved. The lack of optimization techniques for nonmonotonic DL reasoning is a further obstacle to the practical application of these logics.

Given the above motivations (more extensively articulated in the full paper), in this work we have investigated a new family of nonmonotonic DLs, called \mathcal{DL}^N , aimed at addressing the above drawbacks. The comparison of \mathcal{DL}^N with the major nonmonotonic DLs is summarized in Table 1. Preliminary performance tests have been reported in the full paper and have been later extended in [Bonatti *et al.*, 2015b] using additional optimization techniques that show unparalleled scalability properties over large nonmonotonic knowledge bases, with more than 10^5 axioms. For instance, sub-

*This is an extended abstract of the paper [Bonatti *et al.*, 2015a], integrated with the developments in [Bonatti and Sauro, 2017], conditionally accepted for publication on the AIJ. This research is currently being funded by the European Unions Horizon 2020 research and innovation programme under grant agreement N. 731601.

Table 1: Summary of comparisons with other nonmonotonic DLs

Features	CIRC	DEF	AEL	TYP	RAT	PR	\mathcal{DL}^N
no inheritance blocking	✓	✓	✓			✓	✓
no CWA effects		✓	✓		✓	✓	✓
fine-grained control on role ranges				✓ ⁽¹⁾			✓
detects inconsistent prototypes				(2)		✓ ⁽³⁾	✓
unique deductive closure	✓			✓			✓
preserves tractability					(5)	(5)	✓ ⁽⁴⁾
implicit specificity				✓	✓	✓	✓
other priorities	✓	✓					✓

CIRC, DEF, AEL, TYP, RAT, PR stand, respectively, for Circumscribed DLs [Bonatti *et al.*, 2009; 2011], Default DLs [Baader and Hollunder, 1995a; 1995b], Autoepistemic DLs [Donini *et al.*, 2002], DLs with Typicality [Giordano *et al.*, 2009; 2013], DLs with Rational Closure [Casini and Straccia, 2010; Casini *et al.*, 2013] and their combination with inheritance networks [Casini and Straccia, 2013], and Probabilistic DLs [Lukasiewicz, 2008].

(1) Context dependent, see Example 29 in the full paper.

(2) Only direct conflicts, as in Example 10 in the full paper.

(3) Inconsistency may propagate to the entire KB.

(4) Subsumption and assertion checking only, which suffice for application examples.

(5) Currently proved for \mathcal{EL} (Giovanni Casini, personal communication).

sumption/assertion query answering over nonmonotonic variants of the Gene Ontology with classical role fillers ranges from 0.25 to 2.46 seconds.

2 The Family \mathcal{DL}^N

Let \mathcal{DL} be any description logic. The language of \mathcal{DL}^N is obtained by adding a new concept name NC for each \mathcal{DL} concept C . The new concepts are called *normality* concepts, although the term *standard* would be more appropriate since \mathcal{DL}^N – unlike typicality logics and rational closure – has not been designed to model normality (as formalized by the KLM postulates, for example).

\mathcal{DL}^N has a *utilitarian* purpose, related to McCarthy’s *conventions* [McCarthy, 1986], aimed at making knowledge bases more manageable, e.g. by reducing their size, improving modularity, and so on. We may want to assert that by default a drug has contraindication X not because this is normally so, but rather for mitigating the effects of potential human errors [Rector, 2004; Stevens *et al.*, 2007]: It is safer to signal more contraindications than missing some; accordingly, with the above default assertion, forgetting to define the contraindications of a new drug cannot result in a missing contraindication. The goal of \mathcal{DL}^N is supporting such a variety of knowledge engineering needs in a scalable way.

A \mathcal{DL}^N *knowledge base* is a disjoint union $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$ where \mathcal{S} is a finite set of \mathcal{DL}^N inclusions and assertions and \mathcal{D} is a finite set of *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. If $\delta = (C \sqsubseteq_n D)$, then $\text{pre}(\delta)$ and $\text{con}(\delta)$ denote C and D , respectively.

The informal meaning of $C \sqsubseteq_n D$ is: “*all standard instances, by default, satisfy $C \sqsubseteq D$, unless stated otherwise*”, that is, unless some higher priority axioms entail that some standard instances satisfy $C \sqcap \neg D$; in that case, $C \sqsubseteq_n D$ is *overridden*. The instances of any concept NE are required to satisfy all the DIs that are not overridden in NE .

The priority relation over DIs is denoted by \prec . \mathcal{DL}^N solves automatically only the conflicts that can be settled using \prec . Any other conflict should be regarded as a representation error (cf. the discussion of silent conflict removal in the introduction) and shall be resolved by the knowledge engineer, typically by adding specific DIs. Unresolved conflicts yield inconsistent normality concepts, that can be detected by queries of the form $NC \sqsubseteq \perp$.

The priority relation is a parameter of \mathcal{DL}^N . The full paper considers both the specificity-based priority of rational closure and a simpler kind of specificity adopted by circumscribed DLs [Bonatti *et al.*, 2015a], namely: $\delta_1 \prec \delta_2$ iff

$$\mathcal{KB} \approx \text{pre}(\delta_1) \sqsubseteq \text{pre}(\delta_2) \text{ and } \mathcal{KB} \not\approx \text{pre}(\delta_2) \sqsubseteq \text{pre}(\delta_1).$$

The expression $\mathcal{KB} \models \alpha$ means that α is a \mathcal{DL}^N *consequence* of \mathcal{KB} . Due to space limitations, we do not report the model-theoretic definition of \models and present only its reduction to classical reasoning [Bonatti *et al.*, 2015a]. For all subsumptions and assertions α , $\mathcal{KB} \models \alpha$ holds iff $\mathcal{KB}^\Sigma \models \alpha$,¹ where Σ is the set of normality concepts explicitly occurring in $\mathcal{KB} \cup \{\alpha\}$, and \mathcal{KB}^Σ is a classical knowledge base obtained as follows (recall that $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$):

¹In this classical translation, normality concepts are treated like new concept names.

\mathcal{DL}^N complexity		
\mathcal{DL} complexity	subsumption and assertion checking	knowledge base and concept consistency
P	P	P
ExpTime	ExpTime	ExpTime
N2ExpTime	$P^{N2ExpTime}$	$P^{N2ExpTime}$

All results hold for specificity and other priority relations in P^C , where C is the complexity of subsumption in \mathcal{DL}

Table 2: Some complexity results

First, for all DIs $\delta \in \mathcal{D}$ and all $NC \in \Sigma$, let:

$$\delta^{NC} = (NC \sqcap \text{pre}(\delta) \sqsubseteq \text{con}(\delta)). \quad (1)$$

The informal meaning of δ^{NC} is: “ NC ’s instances satisfy δ ”.

Then, let $\downarrow_{\prec \delta}$ be the operator that removes from any axiom set $\mathcal{S}' \supseteq \mathcal{S}$ all the δ_0^{NC} that do not have higher priority than δ , preserving the axioms of \mathcal{KB} :

$$\mathcal{S}' \downarrow_{\prec \delta} = \mathcal{S} \cup (\mathcal{S}' \setminus \{\delta_0^{NC} \mid NC \in \Sigma \wedge \delta_0 \not\prec \delta\}).$$

Third, let $\delta_1, \dots, \delta_{|\mathcal{D}|}$ be any *linearization* of (\mathcal{D}, \prec) .²

Finally, let $\mathcal{KB}^\Sigma = \mathcal{KB}_{|\mathcal{D}|}^\Sigma$, where the sequence \mathcal{KB}_i^Σ ($i = 1, 2, \dots, |\mathcal{D}|$) is inductively defined as follows:

$$\mathcal{KB}_0^\Sigma = \mathcal{S} \cup \{NC \sqsubseteq C \mid NC \in \Sigma\} \quad (2)$$

$$\mathcal{KB}_i^\Sigma = \mathcal{KB}_{i-1}^\Sigma \cup \{\delta_i^{NC} \mid NC \in \Sigma, \text{ and}$$

$$\mathcal{KB}_{i-1}^\Sigma \downarrow_{\prec \delta_i} \cup \{\delta_i^{NC}\} \not\models NC \sqsubseteq \perp\}. \quad (3)$$

In other words, the above sequence starts with \mathcal{KB} ’s strong axioms \mathcal{S} , extended with the inclusions $NC \sqsubseteq C$, then processes the DIs δ_i in non-increasing priority order.

If δ_i can be consistently added to C ’s prototype, given all higher priority DIs selected so far (which is verified by checking that $NC \not\sqsubseteq \perp$ in line (3)), then its translation δ_i^{NC} is included in \mathcal{KB}^Σ (i.e. δ_i “enters C ’s prototype”), otherwise δ_i is discarded, and we say that δ_i is *overridden in NC* .

The original semantics is refined in [Bonatti and Sauro, 2017] in order to infer $NC \equiv ND$ when $C \equiv D$. The corresponding iterative construction of \mathcal{KB}^Σ is very similar: it suffices to replace equation (2) with

$$\mathcal{KB}_0^\Sigma = \mathcal{S} \cup \{NC \sqsubseteq C \mid NC \in \Sigma\} \cup \{NC \equiv ND \mid \mathcal{S} \models C \equiv D\}.$$

Using the construction of \mathcal{KB}^Σ and the equivalence between $\mathcal{KB} \approx \alpha$ and $\mathcal{KB}^\Sigma \models \alpha$ it is not hard to prove that:

1. If \mathcal{DL} belongs to a deterministic time complexity class C , then \mathcal{DL}^N preserves its complexity;
2. If C is nondeterministic, then \mathcal{DL}^N belongs to P^C .

In particular, \mathcal{DL}^N preserves tractability. Some examples – applying to \mathcal{EL}/DI -lite, \mathcal{ALC} , and \mathcal{SROIQ} , respectively – are reported in Table 2.

²That is, $\{\delta_1, \dots, \delta_{|\mathcal{D}|}\} = \mathcal{D}$ and for all $i, j = 1, \dots, |\mathcal{D}|$, if $\delta_i \prec \delta_j$ then $i < j$. All linearizations yield the same \mathcal{KB}^Σ .

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

Table 3: The \mathcal{KB} formalizing the example

3 Behavior of \mathcal{DL}^N

A DI $\delta = (C \sqsubseteq_n D)$ is roughly similar to a set of defaults stating, for each normality concept NE , that the instances of NE satisfy $C \sqsubseteq D$ unless stated otherwise (equivalently, δ^{NE} holds by default). The full paper and [Bonatti *et al.*, 2015b; Bonatti and Sauro, 2017] contain several examples inspired by real applications, but here – due to space limitations – we resort to an old-fashioned 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 3. 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*:

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

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 the set of such exceptional individuals to be minimized. So, in the direct equivalents of the above knowledge base, *Dolphins* would be inconsistent ($\text{Dolphins} \sqsubseteq \perp$). 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 $\text{Whales} \equiv \{\text{Moby}\}$ while \mathcal{DL}^N would not. ■

(A) $\frac{\mathcal{KB} \approx \alpha \quad \mathcal{KB} \approx \gamma}{\mathcal{KB} \cup \{\alpha\} \approx \gamma}$	α, γ are inclusions/assertions
(B) $\frac{C \sqsubseteq_n D \quad C \sqsubseteq_n E}{C \sqcap D \sqsubseteq_n E}$	(C) $\frac{NC \sqsubseteq D \quad NC \sqsubseteq E}{N(C \sqcap D) \sqsubseteq E}$

 Table 4: Variants of cautious monotony in \mathcal{DL}^N

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; Table 4 illustrates the three variants investigated in [Bonatti and Sauro, 2017]³ through the corresponding versions of cautious monotony.

With respect to version (A) of the postulates, \mathcal{DL}^N is cumulative. Moreover, *if each consistent DL concept has at least one standard instance*, then \mathcal{DL}^N satisfies all postulates. Interestingly, 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]. The cumulativity property shows that DIs significantly differ from default rules, since default logic is not cumulative.

Version (B) is affected by overriding. For instance, if the second premise of this variant of cautious monotony were overridden, then the conclusion would not be supported. On the other hand, if the premises are not overridden, then all the postulates of type (B) are valid in \mathcal{DL}^N . A first interesting remark is that a similar behavior 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).

Version (C) of the postulates is analogous to the postulates satisfied by typicality logics. It can be shown that typicality DLs satisfy these postulates 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 – consequently – does not universally satisfy version (C) of Cut and Cautious Monotony. The latter 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:

$$D \sqsubseteq \neg E \quad (4)$$

$$C \sqsubseteq_n D \quad (5)$$

$$C \sqsubseteq_n E. \quad (6)$$

Here the defeasible inclusions (5) and (6) together contradict the strong inclusion (4), but they have the same priority,

³Version (B) is already studied in the full paper.

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 (4), so (6) is overridden and $N(C \sqcap D)$ is consistent. This shows that version (C) of cautious monotony 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 (4) should not override (6) 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 version (C) hold but cautious monotony (cf. Example 2); if there are no unresolved conflicts, then cautious monotony holds, too. So \mathcal{DL}^N is almost rational also with respect to version (C).

Given these results, it is not surprising that \mathcal{DL}^N returns the standard expected conclusions in many artificial examples occurring in the literature. Over a set of examples due to Sandewall, \mathcal{DL}^N and the DL introduced in [Casini and Straccia, 2013] (that combines rational closure and inheritance networks) make the same inferences if the same priority relation is used, with the exception of the examples involving unresolved conflicts. In these cases, of course, \mathcal{DL}^N highlights the conflicts while the other logic silently removes them (cf. Section 6.5 of the full paper).

4 Further Properties and Results

The following contents of the full paper could not be described in this abstract due to space limitations.

Usage and representation methodologies. We show how to deal with a variety of representation needs, introducing some nonmonotonic design patterns for fine-grained control on overriding and role ranges. The latter address a characteristic property of \mathcal{DL}^N not discussed here (but reported in Table 1), namely, the ability to specify whether role filler may include both standard and nonstandard individuals, or whether roles should range on standard individuals only.

Test case generation and validation. Benchmarking techniques and tools for nonmonotonic DL reasoners are still in a pretty early stage. We propose some general criteria for guiding the generation of benchmark tests, and some ex-post validation metrics to assess the quality of the synthetic test cases. These principles apply to all nonmonotonic DLs.

Inferring DIs. The model theoretic semantics of \mathcal{DL}^N provides a natural way of *inferring* further DIs from those explicitly contained in a KB. This feature might turn out to be helpful in optimizing inferences, e.g. by identifying redundant DIs or restructuring DIs in a way that speeds up automated reasoning. We prove that DI inference is in general more expensive than assertion/subsumption inference; its complexity in the general case is still unknown.

The work on implementations and optimizations is currently being completed with refined modularization, incrementality, and parallelization techniques. Additional experiments with knowledge bases of increasing size are being carried out to complete the assessment of the scalability of \mathcal{DL}^N .

References

- [Baader and Hollunder, 1995a] Franz Baader and Bernhard Hollunder. Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reasoning*, 14(1):149–180, 1995.
- [Baader and Hollunder, 1995b] Franz Baader and Bernhard Hollunder. Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *J. Autom. Reasoning*, 15(1):41–68, 1995.
- [Bonatti and Sauro, 2017] Piero A. Bonatti and Luigi Sauro. On the logical properties of the nonmonotonic description logic \mathcal{DL}^N . *Artif. Intell.* Accepted for publication subject to minor revisions, 2017.
- [Bonatti et al., 2009] Piero A. Bonatti, Carsten Lutz, and Frank Wolter. The complexity of circumscription in DLs. *J. Artif. Intell. Res. (JAIR)*, 35:717–773, 2009.
- [Bonatti et al., 2011] Piero A. Bonatti, Marco Faella, and Luigi Sauro. Defeasible inclusions in low-complexity DLs. *J. Artif. Intell. Res. (JAIR)*, 42:719–764, 2011.
- [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*, volume 9366 of *Lecture Notes in Computer Science*, pages 356–372. Springer, 2015.
- [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, 2013] Giovanni Casini and Umberto Straccia. Defeasible inheritance-based description logics. *J. Artif. Intell. Res. (JAIR)*, 48:415–473, 2013.
- [Casini et al., 2013] Giovanni Casini, Thomas Meyer, Kodylan Moodley, and Ivan José Varzinczak. Towards practical defeasible reasoning for description logics. In Thomas Eiter, Birte Glimm, Yevgeny Kazakov, and Markus Krötzsch, editors, *Description Logics*, volume 1014 of *CEUR Workshop Proceedings*, pages 587–599. CEUR-WS.org, 2013.
- [Donini et al., 2002] Francesco M. Donini, Daniele Nardi, and Riccardo Rosati. Description logics of minimal knowledge and negation as failure. *ACM Trans. Comput. Log.*, 3(2):177–225, 2002.
- [Giordano et al., 2009] Laura Giordano, Nicola Olivetti, Valentina Gliozzi, and Gian Luca Pozzato. ALC + T: a preferential extension of description logics. *Fundam. Inform.*, 96(3):341–372, 2009.
- [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.
- [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, 1995] Daniel J. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
- [Lukasiewicz, 2008] Thomas Lukasiewicz. Expressive probabilistic description logics. *Artif. Intell.*, 172(6-7):852–883, 2008.
- [McCarthy, 1986] John McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.
- [Rector, 2004] Alan L. Rector. Defaults, context, and knowledge: Alternatives for OWL-indexed knowledge bases. In *Pacific Symposium on Biocomputing*, pages 226–237. World Scientific, 2004.
- [Stevens et al., 2007] Robert Stevens, Mikel Egaña Aranguren, Katy Wolstencroft, Ulrike Sattler, Nick Drummond, Matthew Horridge, and Alan L. Rector. Using OWL to model biological knowledge. *International Journal of Man-Machine Studies*, 65(7):583–594, 2007.