Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Journal Track

Optimizing the Computation of Overriding in DLN (Extended Abstract)*

Piero A. Bonatti[1] , **Iliana Petrova**[2] and **Luigi Sauro**[1]
[1]Università di Napoli Federico II
[2]Inria, Sophia Antipolis, France
{pieroandrea.bonatti, luigi.sauro}@unina.it, iliana.petrova@inria.fr

The formulation of knowledge typically involves the use of by-default axioms that can be overridden in order to accommodate specific exceptions. Biologists, in particular, have commonly followed an incremental approach to introduce exceptions to general properties. For instance, in humans, the heart is typically positioned on the left side of the thorax. However, there are individuals born with a condition called situs inversus, where the heart is located on the opposite side of the body. Similarly, eukaryotic cells are generally defined as having a nucleus, but mammalian red blood cells, in their mature stage, do not possess a nucleus.

Similarly, privacy policies generally include default conditions, such as *open* and *closed* policies, conflict resolution methods such as *denials take precedence*, and authorization inheritance with exceptions [Bonatti and Samarati, 2003].

Description logics (DLs), which underlie the Semantic Web standard OWL2, do not allow to express defeasible knowledge and exceptions. Consequently, several authors proposed different nonmonotonic extentions as a useful means to address this limitation [Bonatti *et al.*, 2011a; Bonatti *et al.*, 2010; Bonatti *et al.*, 2011b; Donini *et al.*, 2002; Giordano *et al.*, 2013; Giordano *et al.*, 2009; Giordano *et al.*, 2015; Casini and Straccia, 2010; Bonatti, 2019].

In this context, $\mathcal{DL}^N$ [Bonatti *et al.*, 2015a; Bonatti and Sauro, 2017; Bonatti *et al.*, 2015b] is a recent family of nonmonotonic DL which derives from a utilitarian approach to nonmonotonic logic. The primary objective of this approach is to meet the practical requirements of ontology designers, which are highlighted by a number of instances in the literature on biomedical ontologies and semantic web policies. $\mathcal{DL}^N$ features *normality concepts* $\mathsf{N}C$ to denote the standard/prototypical instances of a concept C, and prioritized *defeasible inclusions* (DIs) $C \sqsubseteq_n D$ that mean (roughly speaking): "*by default, all prototypical instances that satisfy C satisfy also D, unless stated otherwise*", that is, unless some higher priority axioms contradict this implication; in this case, $C \sqsubseteq_n D$ is *overridden*. The standard/prototypical instances of C are required to satisfy all the DIs that are not overridden in C.

$\mathcal{DL}^N$ adopts the simplest possible criterion for overriding, that is, inconsistency with higher priority axioms. Conflicts between DIs that cannot be resolved with priorities are regarded as knowledge representation errors and are to be fixed by the knowledge engineer (typically, by adding specific DIs). As a consequence, all the normal instances of a concept C conform to the same set of default properties, sometimes called *prototype*. Here is a summary of the main strenghts of $\mathcal{DL}^N$.

No inheritance blocking. Most of the logics grounded on preferential semantics and rational closure block the inheritance of *all* default properties towards exceptional subclasses (as opposed to overriding only the properties that are modified in those subclasses). $\mathcal{DL}^N$'s overriding mechanism does not suffer from this drawback.

No undesired CWA effects. Many nonmonotonic DLs extend default properties to as many individuals as possible, thereby introducing CWA (i.e. closed-world assumption) effects that clash with the intended behavior of ontologies. $\mathcal{DL}^N$ does not introduce any CWA effect because it does not force individuals to be normal, unless explicitly stated otherwise.

Control on priorities. Since priorities are not fixed a priori in $\mathcal{DL}^N$, knowledge engineers can adapt them to their needs. In principle, it is possible to override DIs based on temporal criteria (which may be useful in legal ontologies and ontology versioning), define default conflict resolution criteria, and even use rational closure's specificity-based axiom ranking. The logics derived from inheritance networks, preferential semantics, and rational closure can only support their fixed, specificity-based overriding criterion.

Default role fillers. $\mathcal{DL}^N$ axioms can specify whether a role should range only over normal individuals or not. Some logics are completely unable to apply default properties to role values.[1] Some others cannot switch this inference off when it is not desired. Only $\mathcal{DL}^N$ and $\mathcal{ALC} + \mathbf{T}_{\min}$ make it possible to control this kind of inference.

Inconsistent prototype detection. $\mathcal{DL}^N$ facilitates the identification of all conflicts that cannot be resolved with priorities (via consistency checks over normality concepts), because their correct resolution is application dependent and should require human intervention.

* This is an extended abstract of [Bonatti *et al.*, 2022].

[1]This is the case for rational closure. Recently, in [Pensel, 2019], a solution has been proposed for $\mathcal{EL}$ with $\bot$. It is unclear how to extend it to more expressive DLs, and it is not possible to "turn off" the application of default rules to role fillers.

Unique deductive closure. As a result of automated conflict resolution, several nonmonotonic logics yield multiple deductive closures, corresponding to all the alternative ways of solving each conflict. $\mathcal{DL}^N$ is one of the logics that has a unique closure.

Generality. $\mathcal{DL}^N$ can be uniformly applied to all description logics up to the standard OWL2-DL (i.e. the logic $\mathcal{SROIQ}(\mathcal{D})$). Typicality logics and rational closure, instead, are limited to logics that satisfy the *disjoint union model property*. Recently, it has been shown that for expressive DLs that do not enjoy this property, syntactic inference does not match semantics [Bonatti, 2019]. The same paper introduces *stable rational closure* that solves the generality problem for rational closure, but re-introduces the issue of multiple (or non existent) deductive closures. It is currently not clear how to design a logic that satisfies the KLM postulates, is fully general, and yields a unique closure for all knowledge bases.

Low complexity. $\mathcal{DL}^N$ preserves the tractability of subsumption and instance checking for all low-complexity DLs, including the rich tractable logics $\mathcal{EL}^{++}$ and DL-lite$^{(\mathcal{HN})}_{Horn}$. Currently, no other nonmonotonic DL enjoys this property to the same extent. Rational closure has been proved to be tractable for $\mathcal{EL}$ extended with $\perp$ [Casini *et al.*, 2019; Pensel, 2019]. Some logics, such as [Casini and Straccia, 2010; Casini *et al.*, 2013; Casini and Straccia, 2013; Giordano *et al.*, 2009; Giordano *et al.*, 2013], preserve the asymptotic complexity of ExpTime-complete DLs like $\mathcal{ALC}$. More generally, $\mathcal{DL}^N$ preserves the asymptotic complexity of all the DLs that belong to a deterministic complexity class that contains P. For nondeterministic complexity classes $\mathcal{C}$, an upper bound is $P^{\mathcal{C}}$.

In [Bonatti *et al.*, 2015a; Bonatti and Sauro, 2017] the semantic properties of $\mathcal{DL}^N$ and the computational complexity of the related reasoning tasks have been thoroughly studied. As mentioned above $\mathcal{DL}^N$ preserves the tractability of low-complexity DLs, this opens the way to processing very large nonmonotonic KBs within these fragments. For practical purposes, however, asymptotic tractability alone is insufficient. $\mathcal{DL}^N$ reasoning is based on an iterative procedure that, given the signature of the queries of interest, discards overridden inclusions and transforms the other defeasible inclusions into classical axioms. In the worst case the number of concept consistency checks carried out by this reduction is quadratic in both the knowledge base and of input signature size. These consistency checks are conducted on different subsets of the knowledge base, which are generally uncomparable. This is why they cannot be computed by a single classification of the knowledge base.

In [Bonatti *et al.*, 2015a], a preliminary implementation of a DLN reasoner has clearly demonstrated that such a quadratic dependence can significantly slow down the execution of the computation, even when the engine takes advantage of the incremental reasoning facilities native to state-of-the-art reasoners like ELK. Consequently, practical reasoning on large knowledge bases – such as biomedical ontologies – requires ad hoc optimizations to detect and prune unnecessary computations during the reduction to classical DL. For this purpose, we propose two optimization techniques that effectively speed up reasoning.

The first optimization focuses on discarding irrelevant axioms for a given query by adapting a classical module extraction algorithm to DLN. This technique also reduces the number of iterations required for the reduction to classical DL. However, adapting classical module extractors to DLN is a challenging task due to the nonmonotonic nature of its inferences.

As mentioned earlier, the initial DLN reasoner utilized the incremental reasoning mechanisms of the underlying classical reasoner. The second optimization, known as the optimistic method, aims to minimize the number of retractions, which are typically the most computationally expensive operations in incremental reasoning.

The efficiency of each optimization and their combined effect is evaluated through experimental analysis. To ensure realistic test cases, we conducted experiments on nonmonotonic versions of prominent biomedical ontologies, including Gene Ontology, Fly Anatomy, and SNOMED. These ontologies not only find applications in specific scenarios but are also widely adopted as benchmarks for assessing performance. To the best of our knowledge, this study represents the first exploration in the field of Description Logics to apply nonmonotonic reasoning techniques to knowledge bases containing a significant number of axioms, ranging between approximately $\sim 20,000$ and $\sim 30,000$.

First, we considered the nonmonotonic versions of Fly Anatomy (>19K axioms) and Gene Ontology (>28K axioms) obtained by transforming classical inclusions into defeasible inclusions. In all of these experiments the average query response time is below one second, and mostly below 0.5 seconds. Similar experiments based on SNOMED (>290K axioms) yield response times between 1.5 and 4.6 seconds, that are compatible with a wide range of use cases. In another set of experiments, we have injected random DIs in Fly Anatomy, Gene Ontology, and SNOMED, thereby increasing their size up to 25%, and introducing random dependencies between different parts of the knowledge base. For the smaller ontologies Fly Anatomy and Gene Ontology, response times are always below 1.5 minutes, and less than 4 seconds in most test cases. For SNOMED, the hardest test cases ($\sim$ 72K additional axioms, 250 normality concepts in the knowledge base) took approximately 1 hour and 45 minutes, while the simplest test cases (N-free) are completed within approximately 1 minute.

In summary, the results obtained from DLN ontologies containing up to 35K axioms demonstrate that the combination of both optimizations significantly reduces computation time by up to four orders of magnitude. This improvement allows for the processing of subsumption queries within response times suitable for various practical applications, including interactive query answering. For knowledge bases derived from SNOMED, which are approximately ten times larger, response times may vary depending on several crucial performance factors. These factors include the number of normality concepts present in the knowledge bases and the extent of logical dependencies between different concepts. In test cases where the size and structure of SNOMED are preserved, the response times remain compatible with interactive query answering. However, in fully random test cases, the response times range from 11 seconds to nearly two hours.

The challenge of developing accurate module extraction techniques for nonmonotonic logics remains an open problem in general. Even within the realm of $\mathcal{DL}^N$, our proofs do not encompass module extraction methods based on principles other than locality. Additionally, the potential applications of module extraction extend beyond optimization. Conducting a comprehensive investigation into module extraction for nonmonotonic reasoning represents an interesting subject for future research.

References

[Bonatti and Samarati, 2003] Piero A. Bonatti and Pierangela Samarati. Logics for authorization and security. In *Logics for Emerging Applications of Databases*, pages 277–323. Springer, 2003.

[Bonatti and Sauro, 2017] Piero A. Bonatti and Luigi Sauro. On the logical properties of the nonmonotonic description logic dl^{n}. *Artif. Intell.*, 248:85–111, 2017.

[Bonatti *et al.*, 2010] Piero A. Bonatti, Marco Faella, and Luigi Sauro. EL with default attributes and overriding. In *Int. Semantic Web Conf. (ISWC 2010)*, volume 6496 of *LNCS*, pages 64–79. Springer, 2010.

[Bonatti *et al.*, 2011a] Piero A. Bonatti, Marco Faella, and Luigi Sauro. Adding default attributes to EL++. In Wolfram Burgard and Dan Roth, editors, *AAAI*. AAAI Press, 2011.

[Bonatti *et al.*, 2011b] 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*, pages 356–372, 2015.

[Bonatti *et al.*, 2022] Piero Andrea Bonatti, Iliana Mineva Petrova, and Luigi Sauro. Optimizing the computation of overriding in DLN. *Artificial Intelligence*, 311:103764, 2022.

[Bonatti, 2019] Piero A. Bonatti. Rational closure for all description logics. *Artif. Intell.*, 274:197–223, 2019.

[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.

[Casini *et al.*, 2019] Giovanni Casini, Umberto Straccia, and Thomas Meyer. A polynomial time subsumption algorithm for nominal safe $\mathrm{elo}^{\perp}$ under rational closure. *Inf. Sci.*, 501:588–620, 2019.

[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.

[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.

[Pensel, 2019] Maximilian Pensel. *A Lightweight Defeasible Description Logic in Depth*. PhD thesis, TU Dresden, 2019.