Minimal Model Reasoning in Description Logics: Don't Try This at Home! (Extended Abstract)

Federica Di Stefano[1], Quentin Manière[2,3], Magdalena Ortiz[1] and Mantas Šimkus[1]

[1]Institute of Logic and Computation, TU Wien
[2]Department of Computer Science, Leipzig University, Germany
[3]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, Germany

Abstract

We summarize our recent work [1] on minimal model reasoning in lightweight description logics.

Keywords

Minimal model reasoning, Description logics, Circumscription, Complexity of reasoning

1. Introduction

Reasoning with *minimal models* has always been at the core of many non-monotonic formalisms, such as *default logic* [2], *circumscription* [3], or *answer set programming* [4]. Despite it capturing the attention of the KR community over the years, there are still big gaps in our understanding of minimal model reasoning in *Description Logics (DLs)*. When reasoning from a knowledge base, *minimal models* provide a natural and intuitive counterpart to traditional open-world semantics and classical entailment, which can easily *exclude* some expected consequences (e.g., a query may be not entailed due to a counter-example model that includes unexpected and unjustified facts). Consider the assertions

$$\mathsf{ScandCountry}(no), \mathsf{ScandCountry}(se), \mathsf{ScandCountry}(dk),$$
$$\mathsf{NatoMember}(no), \mathsf{NatoMember}(se), \mathsf{NatoMember}(dk)$$

Under the classical semantics, the inclusion $\mathsf{ScandCountry} \sqsubseteq \mathsf{NatoMember}$ is not entailed by the above assertions, since there may be unknown Scandinavian countries that are not in NATO. In contrast, considering only those models in which all facts are strictly necessary and justified may lead to more intuitive reasoning, i.e., every Scandinavian country is in NATO.

Predicate minimization has been explored in the context of *circumscribed DLs* , but most existing results spell out the high complexity that results from combining minimized predicates with varying or fixed predicates; see, e.g., [5, 6]. Specifically, when minimized roles and varying predicates are allowed, reasoning becomes quickly undecidable. Except for sporadic results [7, 8], the case of purely *minimal models*, where nothing can be removed from the extension of any predicate while preserving modelhood, remained largely unexplored.

The present extended abstract summarizes our recent work [1]. In [1], we investigate the complexity of reasoning in lightweight DLs in the $\mathcal{EL}$ and DL-Lite family under the *minimal model semantics*, providing the following contributions.

• We show that concept satisfiability in a minimal model is undecidable for the DL $\mathcal{EL}$. The decidability status of minimal model reasoning has been open for several years, and the negative outcome is somewhat surprising. Since the reduction does not use the $\top$-concept, the result carries over a restricted class of *guarded tuple generating dependencies (TGDs)*.

DL 2025: 38th International Workshop on Description Logics, September 3–6, 2025, Opole, Poland
federica.stefano@tuwien.ac.at (F. Di Stefano); quentin.maniere@uni-leipzig.de (Q. Manière); magdalena.ortiz@tuwien.ac.at (M. Ortiz); mantas.simkus@tuwien.ac.at (M. Šimkus)
0000-0002-6570-7598 (F. Di Stefano); 0000-0001-9618-8359 (Q. Manière); 0000-0002-2344-9658 (M. Ortiz); 0000-0003-0632-0294 (M. Šimkus)

• We show that decidability can be regained by imposing two simple *acyclicity conditions* on the TBoxes, namely *strong acyclicity* [7] and *weak acyclicity* [9, 10, 11]. We show that concept satisfiability in minimal models not only becomes decidable, but it is NExp - complete in strongly-acyclic $\mathcal{ELIO}_\bot$, and $\mathsf{NExp}^{\mathsf{NP}}$-complete in weakly-acyclic $\mathcal{ELIO}_\bot$. Furthermore, for the weakly-acyclic $\mathcal{ELIO}_\bot$ we show that concept satisfiability is Σ_2^P-complete in data complexity. Remarkably, our lower bounds hold already for $\mathcal{EL}$.
• We conclude the paper with a minor excursion into DL-Lite, showing that concept satisfiability in minimal models is already ExpSpace-hard for DL-Lite$_{\text{horn}}$.

2. Minimal Model Semantics and Contributions

We refer to [12] for preliminaries on the DLs studied in this paper. We remark that, unless stated otherwise, we make the unique name assumption (UNA).

Definition 1. *Given two interpretations $\mathcal{I}$ and $\mathcal{J}$, we let $\mathcal{I} \subseteq \mathcal{J}$ if*

(i) $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}}$ *and* $a^{\mathcal{I}} = a^{\mathcal{J}}$*, for all* $a \in N_I$*;*
(ii) $p^{\mathcal{I}} \subseteq p^{\mathcal{J}}$*, for all predicates* $p \in N_C \cup N_R$*.*

We write $\mathcal{I} \subsetneq \mathcal{J}$ if $\mathcal{I} \subseteq \mathcal{J}$ and $p^{\mathcal{I}} \subsetneq p^{\mathcal{J}}$ for some $p \in N_C \cup N_R$. We call $\mathcal{I}$ a minimal model *of a KB $\mathcal{K}$, if (a) $\mathcal{I} \models \mathcal{K}$, and (b) there exists no $\mathcal{J} \subsetneq \mathcal{I}$ such that $\mathcal{J} \models \mathcal{K}$.*

Observe that the relation $\subsetneq$ coincides with the preference relation induced by a circumscription pattern where all predicates are minimized [5]. The reasoning task that we focus on is *concept satisfiability in a minimal model* (MinModSat for short) defined as follows: Given an $\mathcal{L}$ KB $\mathcal{K}$ and an $\mathcal{L}$ concept C, decide whether there exists a minimal model $\mathcal{I}$ of $\mathcal{K}$ with $C^{\mathcal{I}} \neq \varnothing$. Other reasoning tasks are outside the scope of this work. We remark that traditional reductions between basic reasoning tasks do not directly apply to minimal model reasoning.

Example 1. *Take a TBox* $\mathcal{T} = \{\mathsf{Fan} \sqsubseteq \exists \mathsf{likes}.\mathsf{Movie} \qquad \mathsf{Critic} \sqsubseteq \exists \mathsf{dislikes}.\top\}$ *stating that (movie)* fans *must* like *some movie, while* critics *always* dislike *something. Consider also ABoxes* $\mathcal{A}_1 = \{\mathsf{Fan}(ann)\}$ *and* $\mathcal{A}_2 = \{\mathsf{Fan}(ann), \mathsf{Critic}(bob)\}$*. We are interested in the satisfiability of the concept* $\mathsf{Movie} \sqcap \exists \mathsf{dislikes}^-.\top$*, i.e., the existence of a movie that is disliked by someone. Observe that C is not satisfiable in a minimal model of $\mathcal{K}_1 = (\mathcal{T}, \mathcal{A}_1)$, because $\mathcal{K}_1$ has no justification of an object (person) that dislikes something. However, the concept is satisfiable in a minimal model of $\mathcal{K}_2 = (\mathcal{T}, \mathcal{A}_2)$ (in this model ann likes a movie that bob dislikes).*

Undecidability. We now state our first and most surprising major result: minimal model reasoning is undecidable already in $\mathcal{EL}$.

Theorem 1. *MinModSat in $\mathcal{EL}$ is undecidable. This holds even if the $\top$-concept is disallowed.*

This result, as well as further complexity lower bounds, heavily relies on the *flooding* technique. Known as *saturation* in disjunctive logic programming [13], this technique simulates the universal quantification required for minimization, *i.e.*, testing that *all* substructures are *not models*. Intuitively, a "flooded" interpretation contains objects that satisfy a given disjunctive concept in more than one way. At the core of this are *cyclic dependencies* between some concept names A_1, A_2 that may appear together in some disjunction $A_1 \sqcup A_2$ on the right-hand-side of a concept inclusion. Intuitively, verifying that $e \in (A_1 \sqcap A_2)^{\mathcal{I}}$ holds in a minimal model $\mathcal{I}$ may require a *case analysis*: checking that $e \in A_1^{\mathcal{I}}$ implies $e \in A_2^{\mathcal{I}}$, and that $e \in A_2^{\mathcal{I}}$ implies $e \in A_1^{\mathcal{I}}$. Such case-based verification can be used for testing for crucial properties (errors in a coloring, in a grid construction, etc.), and a flooded minimal model implies that every possible way of avoiding the flooding failed, thus implicitly quantifying over the domain of the structure. As $\mathcal{EL}$ concepts do not support disjunctions, another key-ingredient in our proof is the simulation of those. This is achieved by forcing (via minimality) the role successor of an element to point to *some* individual, and then read-off *which one* was chosen using existentially qualified concepts.

In our undecidability proof, we rely heavily on cyclic inclusions, and it is thus natural to turn our attention to *acyclic TBoxes*, for which minimal model reasoning becomes more manageable.

Strong Acyclicity. Following [7], we define the *dependency graph* $DG(\mathcal{T})$ of an $\mathcal{ALCIO}$ TBox $\mathcal{T}$ and say that $\mathcal{T}$ is *strongly-acyclic* if $DG(\mathcal{T})$ is acyclic and no node is reachable from $\top$. This notion can be seen as a generalization of the one usually considered for terminologies (e.g., in [5]), which is satisfied, for example, by the well-known medical terminology SNOMEDCT. To obtain decidability of MinSat in strongly-acyclic KBs, we rely on results on *pointwise circumscription* [8], where minimization is allowed only *locally*, at one domain element, in contrast to our definition of minimal models, in which predicates are minimized *globally*, across the entire interpretation. Notably, we inherit an NExp complexity upper bound for strongly acyclic KBs in $\mathcal{ALCIO}_{d\leq 1}$, which is the fragment of $\mathcal{ALCIO}$ with modal depth one [8], as minimal models and pointwise minimal models coincide [7]. The results also holds for $\mathcal{ELIO}_\bot$, as it can be reduced to $\mathcal{ALCIO}_{d\leq 1}$ using standard normalization techniques.

Theorem 2. *MINMODSAT in strongly-acyclic $\mathcal{ALCIO}_{d\leq 1}$ and in strongly-acyclic $\mathcal{ELIO}_\bot$ is* NExp*-complete. The lower bound holds already for MINMODSAT in strongly-acyclic $\mathcal{EL}$.*

The proof of the lower bound exploits the following example that illustrates how strongly-acyclic $\mathcal{EL}$ may require exponentially-large models to satisfy a concept of interest.

Example 2. *To generate a binary tree with 2^n leaves, consider the assertion $\mathsf{L}_0(a)$ and axioms $\mathsf{L}_i \sqsubseteq \exists \mathsf{r}_i.\mathsf{L}_{i+1} \sqcap \exists \mathsf{l}_i.\mathsf{L}_{i+1}$ for all $0 \leq i < n$. We want to ensure that all leaves are different objects. For this, we add axioms that attempt to produce a second tree* starting from its leaves. *The latter are identified by concept L'_0, which is made available at leaves of the first tree via $\mathsf{L}_n \sqsubseteq \mathsf{L}'_0$. Further levels of the second tree, towards its root, are generated with the following axioms for $0 \leq j < n$:*

$$\begin{matrix}\mathsf{Left}(o)\\ \mathsf{Right}(o')\end{matrix} \quad \mathsf{L}'_j \sqsubseteq \exists \mathsf{pick}.\top \quad \begin{matrix}\mathsf{L}'_j \sqcap \exists \mathsf{pick}.\mathsf{Left} \sqsubseteq \exists \mathsf{l}'_j.\mathsf{L}'_{j+1,l}\\ \mathsf{L}'_j \sqcap \exists \mathsf{pick}.\mathsf{Right} \sqsubseteq \exists \mathsf{r}'_j.\mathsf{L}'_{j+1,r}\end{matrix} \quad \mathsf{L}'_{j+1,l} \sqcap \mathsf{L}'_{j+1,r} \sqsubseteq \mathsf{L}'_{j+1}$$

A minimal model can only satisfy the concept L'_n if its interpretation of the first tree produces at least 2^n instances of L_n, i.e. of L'_0. Indeed, in a minimial model, each element d in some L'_j has a unique pick*-successor e_d, which, in turn, provides d with either a unique l'_j-successor satisfying $\mathsf{L}'_{j+1,l}$ (if e_d is the interpretation of o), or with a unique r'_j-successor satisfying $\mathsf{L}'_{j+1,r}$ (if e_d is the interpretation of o'). Hence, if there are m elements satisfying L'_{j+1}, then there exist at least $2m$ elements satisfying L'_j. By induction, if L'_n is satisfied, then there are at least 2^n instances of L'_0.*

Weak Acyclicity. We also turn to *weak acyclicity*, which is an important notion for TGDs in the database literature. It relaxes strong acyclicity by annotating some edges in $DG(\mathcal{T})$ as $\star$-edges, intuitively those witnessing axioms of shape $\mathsf{A} \sqsubseteq \exists r.\mathsf{B}$, and defining a TBox $\mathcal{T}$ as *weakly-acyclic* if there is no cycle in $DG(\mathcal{T})$ that goes through a $\star$-edge and no node is reachable from $\top$ in $DG(\mathcal{T})$. We establish a small model property for weakly-acyclic $\mathcal{ELIO}_\bot$, which leads to the following result also considering data complexity.

Theorem 3. *MINMODSAT in weakly-acyclic $\mathcal{ELIO}_\bot$ is $\mathsf{NExp}^{\mathsf{NP}}$-complete. MINMODSAT for weakly acyclic $\mathcal{ELIO}_\bot$ is Σ_2^P-complete in data complexity. Lower bounds hold already for $\mathcal{EL}$.*

3. Perspectives

DL-Lite. We do not study the feasibility of MINMODSAT in the DL-Lite family, but only present one interesting result hinting that the problem will not be easy. In very stark contrast to the previously known NL-membership for MINMODSAT in DL-Lite$_{\text{core}}$[14], already in DL-Lite$_{\text{horn}}$we have ExpSpace-hardness. We hope that this variant and even more expressive extensions like DL-Lite$_{\text{bool}}$ may be decidable, and plan to look for tight matching complexity bounds.

Theorem 4. *MINMODSAT in* DL-Lite$_{\text{horn}}$ *is* ExpSpace*-hard.*

Tuple Generating Dependencies. $\mathcal{EL}$ without $\top$ can be seen as a small fragment of *Tuple Generating Dependencies (TGDs)*, which are prominent in the Database Theory literature (see, e.g., [9, 15]. Thus our lower bounds carry over to minimal model reasoning in TGDs, for problems like *brave entailment* of an atom, or for checking non-emptiness of a relation in some minimal model of a database and input TGDs. Specifically, an $\mathcal{EL}$ TBox without $\top$ can be converted into the so-called *guarded TGDs* with relations of arity at most 2. Minimal model reasoning over TGDs has been explored in [10], where an undecidability result was achieved using relations of arities up to 4 in the context of the *stable model semantics*. Our Theorem 1 implies that checking the existence of a stable model for *normal guarded* TGDs is undecidable already for theories of the form $\Sigma \cup \{\neg g(\vec{t}) \rightarrow \bot\}$, where Σ has negation-free guarded TGDs with relations of arity ≤ 2, and $g(\vec{t})$ is a ground atom. Similarly, our Σ_2^P lower bound in data complexity can be used to improve the Π_2^P lower bound in [10] for weakly acyclic TGDs with stable negation. It remains to be explored whether acyclicity conditions and pointwise minimization might also be useful in the richer setting of TGDs.

Acknowledgments

This work was partially supported by the Austrian Science Fund (FWF) projects PIN8884924, P30873 and 10.55776/COE12.
The authors acknowledge the financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] F. Di Stefano, Q. Manière, M. Ortiz, M. Šimkus, Minimal Model Reasoning in Description Logics: Don't Try This at Home!, (To appear in) Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR), 2025.
[2] R. Reiter, A logic for default reasoning, Artificial intelligence 13 (1980) 81–132. doi:10.1016/0004-3702(80)90014-4.
[3] J. McCarthy, Circumscription: a form of non-monotonic reasoning, Artificial intelligence 13 (1980) 27–39. doi:10.1016/0004-3702(80)90011-9.
[4] M. Gelfond, V. Lifschitz, The stable model semantics for logic programming, in: Proceedings of the 5th International Logic Programming Conference and Symposium (ICLP/SLP), 1988, pp. 1070–1080.
[5] P. A. Bonatti, C. Lutz, F. Wolter, The complexity of circumscription in DLs, Journal of Artificial Intelligence Research 35 (2009) 717–773. doi:10.1613/jair.2763.
[6] C. Lutz, Q. Manière, R. Nolte, Querying Circumscribed Description Logic Knowledge Bases, in: Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR), 2023, pp. 482–491. doi:10.24963/kr.2023/47.
[7] F. Di Stefano, M. Šimkus, Equilibrium Description Logics: Results on Complexity and Relations to Circumscription, in: Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning (KR), 2024, pp. 306–316. doi:10.24963/kr.2024/29.
[8] F. Di Stefano, M. Ortiz, M. Simkus, Description logics with pointwise circumscription, in: Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI), 2023, pp. 3167–3175. doi:10.24963/ijcai.2023/353.

[9] R. Fagin, P. G. Kolaitis, R. J. Miller, L. Popa, Data exchange: semantics and query answering, Theoretical Computer Science 336 (2005) 89–124. doi:10.1016/j.tcs.2004.10.033.

[10] M. Alviano, M. Morak, A. Pieris, Stable model semantics for tuple-generating dependencies revisited, in: Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS), 2017, p. 377–388. doi:10.1145/3034786.3034794.

[11] B. C. Grau, I. Horrocks, M. Krötzsch, C. Kupke, D. Magka, B. Motik, Z. Wang, Acyclicity notions for existential rules and their application to query answering in ontologies, Journal of Artificial Intelligence Research 47 (2013) 741–808. doi:10.1613/jair.3949.

[12] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017. doi:10.1017/9781139025355.

[13] T. Eiter, G. Gottlob, On the computational cost of disjunctive logic programming: Propositional case, Annals of Mathematics and Artificial Intelligence 15 (1995) 289–323. doi:10.1007/BF01536399.

[14] P. A. Bonatti, F. Di Stefano, M. Ortiz, M. Simkus, Circumscription in DL-Lite: Progress Report, in: Description Logics, volume 3515 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.

[15] A. Calì, G. Gottlob, A. Pieris, Towards more expressive ontology languages: The query answering problem, Artificial Intelligence 193 (2012) 87–128. doi:https://doi.org/10.1016/j.artint.2012.08.002.