Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Minimal Model Reasoning in Description Logics: Don't Try This at Home!

Federica Di Stefano[1], **Quentin Manière**[2,3], **Magdalena Ortiz**[1], **Mantas Šimkus**[1]
[1]Institute of Logic and Compuation, TU Wien
[2]Department of Computer Science, Leipzig University, Germany
[3]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, Germany
quentin.maniere@uni-leipzig.de, {federica.stefano|magdalena.ortiz|mantas.simkus}@tuwien.ac.at

Abstract

Reasoning with minimal models has always been at the core of many knowledge representation techniques, but we still have only a limited understanding of this problem in *Description Logics (DLs)*. Minimization of *some* selected predicates—letting the remaining predicates *vary* or be *fixed*, as proposed in *circumscription*—has been explored and exhibits high complexity. The case of 'pure' minimal models, where the extension of *all* predicates must be minimal, has remained largely uncharted. We address this problem in popular DLs and obtain surprisingly negative results: concept satisfiability in minimal models is undecidable already for $\mathcal{EL}$. This undecidability also extends to a very restricted fragment of *tuple-generating dependencies*. To regain decidability, we impose acyclicity conditions on the TBox that bring the worst-case complexity below double exponential time and allow us to establish a connection with the recently studied pointwise circumscription; we also derive results in data complexity. We conclude with a brief excursion to the DL-Lite family, where a positive result was known for DL-Lite$_{\text{core}}$, but our investigation establishes ExpSpace-hardness already for its extension DL-Lite$_{\text{horn}}$.

1 Introduction

Reasoning with *minimal models* has always been at the core of many *Knowledge Representation (KR)* languages. It is most prominent in formalisms for non-monotonic reasoning, from default logic (Reiter 1980) and circumscription (McCarthy 1980) to answer set programming (Gelfond and Lifschitz 1988) and it plays a crucial role in classical KR problems like abduction and diagnosis (Reiter 1987). Finding minimal models and reasoning about them has been a recurring topic in the KR literature for many years; see (Ben-Eliyahu and Dechter 1996; Ben-Eliyahu-Zohary and Palopoli 1997; Lackner and Pfandler 2012; Angiulli et al. 2014; Pfandler, Pichler, and Woltran 2014).

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). In contrast, considering only those models in which all facts are strictly necessary and justified may lead to more intuitive reasoning. We illustrate this in a very simple example.

Example 1. *Under the standard semantics, the inclusion* $\mathsf{ScandCountry} \sqsubseteq \mathsf{NatoMember}$ *is not entailed by the following six assertions, since there may be unknown Scandinavian countries that are not in NATO:* $\mathsf{ScandCountry}(no)$, $\mathsf{ScandCountry}(se)$, $\mathsf{ScandCountry}(dk)$, $\mathsf{NatoMember}(no)$, $\mathsf{NatoMember}(se)$, $\mathsf{NatoMember}(dk)$. *However, the entailment does hold under the minimal model semantics; equivalently, the concept* $\mathsf{ScandCountry} \sqcap \neg\mathsf{NatoMember}$ *is not satisfiable in the minimal models.*

Despite the strong motivation, there are still big gaps in our understanding of minimal model reasoning in *Description Logics (DLs)*. Predicate minimization has been explored in the context of *circumscription in DLs*, but most existing results spell out the high complexity that results from combining minimized predicates with varying or fixed predicates; see, e.g., (Bonatti, Lutz, and Wolter 2009; Lutz, Manière, and Nolte 2023). Specifically, when varying predicates are allowed (e.g. satisfiability of a concept for general circumscription in $\mathcal{ALC}$), reasoning becomes quickly undecidable. But the case of purely *minimal models*, where nothing can be removed from the extension of any predicate while preserving modelhood, remained largely unexplored. It was however established recently that for the DL $\mathcal{ELIO}$—a relatively expressive DL with ExpTime-complete concept satisfiability problem for classical semantics—basic minimal model reasoning becomes undecidable (Di Stefano and Šimkus 2024). A positive result was established for DL-Lite$_{\text{core}}$: here minimal model reasoning exhibits the same worst-case complexity as in the classical case (Bonatti et al. 2023). It is thus natural to explore whether similar positive or negative results can be obtained for other lightweight DLs like $\mathcal{EL}$ or other DL-Lite variants.

In this paper we investigate these questions, and provide 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. It contrasts with the complexity of the classical semantics for $\mathcal{EL}$, which supports tractable reasoning for basic reasoning tasks. Our undecidability proof does not use the $\top$-concept, and it carries over to *guarded tuple generating dependencies (TGDs)* of very restricted shapes.

- To regain decidability, we impose two types of *acyclicity conditions* on the TBox, which are defined in terms of a dependency graph on the predicates of a knowledge base. If we restrict our attention to *strongly acyclic* TBoxes, we can import results from *pointwise circumscription* of (Di Stefano, Ortiz, and Simkus 2023) to show that $\mathcal{ELIO}$ not only becomes decidable, but is feasible in non-deterministic exponential time. We also explore *weak acyclicity*, a common notion in the setting of TGDs in database theory (Fagin et al. 2005; Alviano, Morak, and Pieris 2017; Grau et al. 2013). Weakly acyclic $\mathcal{EL}$ and $\mathcal{ELIO}$ remain decidable; we get tight $\mathsf{NExp}^{\mathsf{NP}}$ bounds on combined complexity, and Σ_2^P on data complexity.
- We conclude the paper with a minor excursion into DL-Lite, but even there we find a challenging panorama: satisfiability in minimal models is already $\mathsf{ExpSpace}$-hard for DL-Lite$_{\text{horn}}$.

An appendix with full proofs can be found in the long version of this paper (Di Stefano et al. 2025).

2 Preliminaries

We briefly recall the syntax and semantics of DLs studied in this paper and refer to (Baader et al. 2017) for more details. We consider countably infinite pairwise disjoint sets $\mathsf{N_C}$, $\mathsf{N_R}$ and $\mathsf{N_I}$ of *concept*, *role* and *individual names*, respectively, and use $\mathsf{N_R^\pm}$ to denote the set $\mathsf{N_R} \cup \{r^- \mid r \in \mathsf{N_R}\}$. *Concepts* in $\mathcal{ALCIO}$ follow the syntax $C := A \mid \{a\} \mid \top \mid \neg C \mid C \sqcap C \mid \exists r.C$, where $A \in \mathsf{N_C}$, $r \in \mathsf{N_R^\pm}$ and $a \in \mathsf{N_I}$. By removing *negated concepts* $\neg C$ from this grammar, we obtain concepts in $\mathcal{ELIO}$; by removing *nominals* $\{a\}$ and requiring $r \in \mathsf{N_R}$ we obtain $\mathcal{ALC}$. The intersection of the $\mathcal{ELIO}$ and $\mathcal{ALC}$ concept languages is called $\mathcal{EL}$. In $\mathcal{ELIO}_\bot$ and $\mathcal{EL}_\bot$ we extend $\mathcal{ELIO}$ and $\mathcal{EL}$ with negation, but only in concepts of the form $\neg\top$, which we equivalently write $\bot$. In $\mathcal{ALCIO}$ we use $C \sqcup D$ as a shortcut for $\neg(\neg C \sqcap \neg D)$ and $\forall r.C$ as a shortcut for $\neg(\exists r.\neg C)$.

Let $\mathcal{L}$ be a DL. A *TBox* $\mathcal{T}$ (in $\mathcal{L}$) is a finite set of *concept inclusions* $C \sqsubseteq D$ where C and D are concepts in $\mathcal{L}$. An *ABox* $\mathcal{A}$ is a finite set of *assertions* of the forms $A(a)$ and $r(a, b)$ with $A \in \mathsf{N_C}$, $r \in \mathsf{N_R^\pm}$ and $a, b \in \mathsf{N_I}$. A pair $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ of a TBox and an ABox is a *knowledge base (KB)*.

The semantics of DLs is defined using interpretations $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty domain and the interpretation function $\cdot^\mathcal{I}$ maps each $A \in \mathsf{N_C}$ to a set $A^\mathcal{I} \subseteq \Delta^\mathcal{I}$, each $r \in \mathsf{N_R}$ to a set of pairs $r^\mathcal{I} \subseteq \Delta^\mathcal{I} \times \Delta^\mathcal{I}$, and each $a \in \mathsf{N_I}$ to an element $a^\mathcal{I} \in \Delta^\mathcal{I}$. The interpretation function extends to all concepts as usual, and we call $\mathcal{I}$ a *model* of a concept C if $C^\mathcal{I} \neq \varnothing$. For α a concept inclusion, assertion, TBox, ABox or KBs, *modelhood* $\mathcal{I} \models \alpha$ is standard. We say that $\mathcal{I}$ makes the *unique name assumption (UNA)* if $a^\mathcal{I} \neq b^\mathcal{I}$ for every $a, b \in \mathsf{N_I}$ with $a \neq b$. When considering $\mathcal{EL}$ and $\mathcal{ELIO}$, we make the UNA unless stated otherwise. In DLs containing $\mathcal{EL}_\bot$ this assumption is irrelevant since the UNA can be simulated in the usual way.

Definition 1. *For 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 $a \in \mathsf{N_I}$;

(ii) $p^\mathcal{I} \subseteq p^\mathcal{J}$ for all predicates $p \in \mathsf{N_C} \cup \mathsf{N_R}$.

We write $\mathcal{I} \subsetneq \mathcal{J}$ if $\mathcal{I} \subseteq \mathcal{J}$ and $p^\mathcal{I} \subsetneq p^\mathcal{J}$ for some $p \in \mathsf{N_C} \cup \mathsf{N_R}$. 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 $\mathcal{J} \models \mathcal{K}$.*

Interpretations with different domains are not comparable according to this definition, which coincides with the preference relation induced by a circumscription pattern where all predicates are minimized (Bonatti, Lutz, and Wolter 2009).

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 $C^\mathcal{I} \neq \varnothing$.

Example 2. *Take a TBox $\mathcal{T}$ stating that (movie)* fans *must* like *some movie, while* critics *always* dislike *something:*

$$\mathsf{Fan} \sqsubseteq \exists \mathsf{likes}.\mathsf{Movie} \qquad \mathsf{Critic} \sqsubseteq \exists \mathsf{dislikes}.\top$$

Consider also ABoxes as follows:

$$\mathcal{A}_1 = \{\mathsf{Fan}(ann)\} \qquad \mathcal{A}_2 = \{\mathsf{Fan}(ann), \mathsf{Critic}(bob)\}$$

We are interested in the satisfiability of the concept $C = \mathsf{Movie} \sqcap \exists \mathsf{dislikes}^-.\top$, 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, C 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).

Since traditional reductions between basic reasoning tasks do not directly apply to minimal model reasoning, we do not study them here and we focus on concept satisfiability only.

3 Undecidability of MinModSat

Before we present our main results, we first provide as a 'warm-up' a proof of Σ_2^P-hardness in data complexity of MinModSat in $\mathcal{ALC}$. The proof is not presented for the complexity result, which is subsumed by tighter bounds in the following sections, but to provide a gentle introduction to the *flooding* technique that will be used heavily in the later reductions. This technique, known as *saturation* in disjunctive logic programming (Eiter and Gottlob 1995), 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*: we may need to check 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.

Example 3. *We show a reduction from (the complement of)* Cert3Col*, a Π_2^P-hard problem (Stewart 1991), to* MinModSat*. As in (Bonatti, Lutz, and Wolter 2009), we define an instance of* Cert3Col *(of size n) as an undirected*

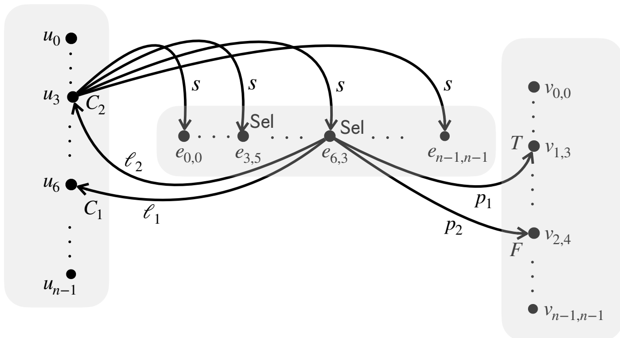

Figure 1: The structure of an ABox $\mathcal{A}$ in Example 3. Roles p_1, p_2 indicate that the first and second literals in the labeling of the edge $(6, 3)$ are $v_{1,3}$ and $v_{2,4}$, respectively. The figure also indicates an interpretation where $v_{1,3}$ and $v_{2,4}$ are respectively true and false, while the vertices 6 and 3 are colored with colors C_1 and C_2. The presence of Sel indicates that the disjunctions of literals at edges $\{3, 5\}$ and $\{6, 3\}$ evaluate to true in the given interpretation.

graph G with vertices $\{0, 1, \ldots, n-1\}$ where each edge is labeled with a disjunction of two literals over the Boolean variables $\{V_{i,j} \mid i, j < n\}$. G is a positive instance of CERT3COL *if for every truth assignment t, the subgraph G_t of G that contains those edges whole labels evaluate to* true *under t is 3-colorable.*

We represent G as an ABox $\mathcal{A}$. We then provide a fixed TBox $\mathcal{T}$ and concept D such that D is satisfiable in a minimal model of $(\mathcal{T}, \mathcal{A})$ iff G is not *a positive instance of* CERT3COL. *The stucture of $\mathcal{A}$ is illustrated in Figure 1.*

Creating the individuals. *To represent the vertices of G, we add the assertions $N(u_0), \ldots, N(u_{n-1})$ to $\mathcal{A}$. Similarly, to represent the propositional variables, we include the assertion $V(v_{i,j})$ for all $0 \leq i, j < n$. To represent the edges of G, we add to $\mathcal{A}$ the assertion $E(e_{i,j})$, for all $0 \leq i, j < n$.*

Connecting the individuals. *We use two roles ℓ_1, ℓ_2 to connect each $e_{i,j}$ to the individuals that represent the vertices i and j: we add $\ell_1(e_{i,j}, u_i)$ and $\ell_2(e_{i,j}, u_j)$ for all $0 \leq i, j < n$. Similarly, using the roles p_1, p_2, we connect each $e_{i,j}$ to the two individuals that represent the literals in the label between i and j: we add $p_1(e_{i,j}, v_{k_1,k_2})$ and $p_2(e_{i,j}, v_{k_3,k_4})$, where v_{k_1,k_2} and v_{k_3,k_4} are the first and second variables in the label of the edge between i and j, respectively. In addition, if v_{k_1,k_2} occurs positively (resp., negatively) in the label, we add $\mathsf{pos}_1(e_{i,j})$ (resp., $\mathsf{neg}_1(e_{i,j})$) to $\mathcal{A}$. The second literal is treated similarly, adding $\mathsf{pos}_2(e_{i,j})$ or $\mathsf{neg}_2(e_{i,j})$ depending on whether v_{k_3,k_4} is positive or negative in the edge represented by $e_{i,j}$. Finally, we use a role s to connect each individual u_k to all the individuals $e_{i,j}$ representing an edge: we add $s(u_k, e_{i,j})$ to $\mathcal{A}$, for all $0 \leq k, i, j < n$. This completes the construction of $\mathcal{A}$.*

Building the TBox. *$\mathcal{T}$ consists of the following inclusions; note that these do not depend on the instance G. We introduce two disjunctions that enforce in each interpretation a 'guess' of a truth value for each individual representing a propositional variable, and a 'guess' of a color assignment for each individual corresponding to a vertex.*

$$V \sqsubseteq T \sqcup F \qquad N \sqsubseteq C_1 \sqcup C_2 \sqcup C_3$$

With the following pair of inclusions we select the edges of G that are labeled with a disjunction that evaluates to true*:*

$$(\mathsf{pos}_i \sqcap \exists p_i.T) \sqcup (\mathsf{neg}_i \sqcap \exists p_i.F) \sqsubseteq \mathsf{Sel} \quad \textit{for } i \in \{1, 2\}$$

Now comes the interesting part: if a selected edge is wrongly colored (i.e. the same color is found at both ends), then all vertices are "flooded" with all colors.

$$\exists s.(\mathsf{Sel} \sqcap \exists \ell_1.C_i \sqcap \exists \ell_2.C_i) \sqsubseteq C_1 \sqcap C_2 \sqcap C_3 \ \textit{ for } i \in \{1, 2, 3\}$$

Finally, we take $C_1 \sqcap C_2$ as our goal concept, which is meant to detect 'flooding' of the structure. It can be easily verified that $(C_1 \sqcap C_2)$ is satisfiable in a minimal model of $(\mathcal{T}, \mathcal{A})$ iff G is not *a positive instance of* CERT3COL.

This reduction will be adapted below to show Σ_2^P-hardness in data complexity for *weakly acyclic* $\mathcal{EL}$ KBs. However, we want to stress that there are no existential concepts on the right-hand-side of inclusions in Example 3, and that all these inclusions can be written in *Disjunctive Datalog (DD)* using rules where all relations have arity ≤ 2. Thus we (slightly) strengthen the Σ_2^P-hardness proof for data complexity of DD in (Eiter, Gottlob, and Mannila 1997), which uses a relation of arity 5.

We now show 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.*

We reduce from RECTTILE, the rectangular tileabilty problem, known to be undecidable (Yang 2014): given a set T of Wang tiles (a.k.a. dominos) and a special color b, decide whether T tiles some finite rectangle with b on its sides.

Consider an instance (T, b) of RECTTILE. We construct an $\mathcal{EL}$ KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ such that $(T, b) \in$ RECTTILE iff the concept Goal is minimally satisfiable w.r.t. $\mathcal{K}$.

Our main challenge is to guarantee that every minimal model satisfying Goal features a rectangular grid representing a tiling of some rectangle. Elements in the grid, further referred to as *nodes*, are identified by the concept Node. We distinguish nine types of positions in the grid, identified using abbreviations of *north*, *south*, *east* and *west*: being one of the four corners (NE, NW, SE, SW), lying on one of the four borders (S, N, E, W), or lying in the *central* part of the grid (C). The following displays an example of how nodes (denoted by their positions) are intended to be arranged:

NW	N	N	N	N	NE
W	C	C	C	C	E
W	C	C	C	C	E
SW	S	S	S	S	SE

We set Pos $:= \{\mathsf{C}, \mathsf{S}, \mathsf{N}, \mathsf{E}, \mathsf{W}, \mathsf{SE}, \mathsf{SW}, \mathsf{NE}, \mathsf{NW}\}$ the set of those nine concept names. We say that p admits p' as a valid horizontal successor, denoted $p \rightsquigarrow_h p'$ if, in the above example, p appears on the same row and more to the left than some p' (*e.g.* $\mathsf{W} \rightsquigarrow_h \mathsf{C}$ and $\mathsf{W} \rightsquigarrow_h \mathsf{E}$, but $\mathsf{W} \not\rightsquigarrow_h \mathsf{SE}$). Similarly, we say that p admits p' as a valid vertical successor, denoted $p \rightsquigarrow_v p'$ if, in the above example, p occurs on the same column and below some p'. Each of these concepts is made available on a dedicated individual in the ABox by adding an assertion $p(a_p)$ for each $p \in \mathsf{Pos}$.

A node e having position p is represented by a connection from e to a_p by the role pos. Each node justifies such a connection using the axiom $\mathsf{Node} \sqsubseteq \exists \mathsf{pos}.\mathsf{Any}$, where Any is an auxiliary concept name that is used in place of $\top$. Note that there is no guarantee that a node connects to one of the individuals a_p. However, such a connection is required to justify further nodes in the grid. Indeed, only the south-west corner is provided as part of the ABox, with assertions $\mathsf{Node}(a)$, $\mathsf{Any}(a)$ and $\mathsf{pos}(a, a_{\mathsf{SW}})$. Further nodes are generated by existing nodes with a valid position and as horizontal or vertical neighbors, which is represented by respective roles h and v. To this end, we add the following axioms:

$$\mathsf{Node} \sqcap \exists \mathsf{pos}.p \sqsubseteq \exists \mathsf{h}.\mathsf{Node} \quad \text{for } p \in \mathsf{Pos} \setminus \{\mathsf{SE}, \mathsf{E}, \mathsf{NE}\}$$
$$\mathsf{Node} \sqcap \exists \mathsf{pos}.p \sqsubseteq \exists \mathsf{v}.\mathsf{Node} \quad \text{for } p \in \mathsf{Pos} \setminus \{\mathsf{NW}, \mathsf{N}, \mathsf{NE}\}.$$

With only these axioms, it should be clear that, in a minimal model, every instance of the Node concept in $\mathcal{I}$ is generated following a path of roles h and v from a and has at most one h-successor and at most one v-successor that is a node. Given an interpretation $\mathcal{I}$, we say that a sequence $d_0, r_1, d_1 \ldots, r_n, d_n$ is an h-v-path of nodes from a to d_n if:

- $d_0 = a^{\mathcal{I}}$ and for every $0 \leq i \leq n$, we have $d_i \in \mathsf{Node}^{\mathcal{I}}$;
- for every $1 \leq i \leq n$, $r_i \in \{\mathsf{h}, \mathsf{v}\}$ and $(d_{i-1}, d_i) \in r_i^{\mathcal{I}}$.

The following claim holds, also in the presence of the remaining inclusions that we add further in this construction:

Claim 1. *Let $\mathcal{I}$ be a minimal model of $\mathcal{K}$. If $d \in \mathsf{Node}^{\mathcal{I}}$, then there exists an* h-v*-path of nodes from a to d. Furthermore, there exists at most one element $e \in \mathsf{Node}^{\mathcal{I}}$ such that $(d, e) \in \mathsf{h}^{\mathcal{I}}$. The same holds for role* v*.*

Using the same mechanism as to 'choose' positions, we ask each node to choose a tile with the following axiom $\mathsf{Node} \sqsubseteq \exists \mathsf{tile}.\mathsf{Any}$ and assertions $t(a_t)$ for each $t \in T$.

We now require the assignments of positions and tiles to be consistent in a minimal model to satisfy respective subgoal-concepts $\mathsf{Subgoal}_1$ and $\mathsf{Subgoal}_2$. More formally, regarding positions, we aim for the following claim:

Claim 2. *Let $\mathcal{I}$ be a minimal model of $\mathcal{K}$ s.t. $a^{\mathcal{I}} \in \mathsf{Subgoal}_1^{\mathcal{I}}$. Then, for every $d \in \mathsf{Node}^{\mathcal{I}}$, there exists a unique $p \in \mathsf{Pos}$, that we denote $pos(d)$, such that $(d, a_p) \in \mathsf{pos}^{\mathcal{I}}$. Furthermore, for every $d, e \in \mathsf{Node}^{\mathcal{I}}$, the following properties hold:*

1. *if $(d, e) \in \mathsf{h}^{\mathcal{I}}$, then $pos(d) \rightsquigarrow_h pos(e)$;*
2. *if $(d, e) \in \mathsf{v}^{\mathcal{I}}$, then $pos(d) \rightsquigarrow_v pos(e)$;*
3. *for every* h-v*-path $d_0, r_1, d_1 \ldots, r_n, d_n$ of nodes from a, there exists (a potentially longer) one $d_0, r_1, d_1 \ldots, r_n, d_n, \ldots, r_{n+k}, d_{n+k}$, with $k \geq 0$, and such that $pos(d_{n+k}) = \mathsf{NE}$;*
4. *if $pos(d) = \mathsf{NE}$, then $d \notin (\exists \mathsf{h})^{\mathcal{I}}$ and $d \notin (\exists \mathsf{v})^{\mathcal{I}}$.*

To achieve the above, we add the rules in Figure 2. Intuitively, we ensure that $\mathsf{Subgoal}_1$ can only be obtained if a concept GoodP is seen at the root a, which is identified with a dedicated assertion $\mathsf{Root}(a)$. To derive GoodP at a given node e, we require all its node successors (there are at most two, due to Claim 1) to already satisfy GoodP and for their respective positions to be valid successors of the position of e. In particular, nodes with position NE trivially satisfy this condition as they expect no successors. Note that this also guarantees Point 3 in the above claim, as nodes along arbitrary long h-v-path of nodes starting from a need not to satisfy GoodP.

We denote H the set of pairs of tiles $(t, t') \in T \times T$ such that the right color of t is the same as the left color of t', so that $(t, t') \in H$ iff t is a valid immediate left-neighbor of t'. Similarly, we denote V the set of pairs of tiles $(t, t') \in T \times T$ such that the top color of t is the same as the bottom color of t'. With rules similar to those on Figure 2, we can ensure consistency of the tiling if the concept $\mathsf{Subgoal}_2$ is satisfied, which is summarized in the following claim:

Claim 3. *Let $\mathcal{I}$ be a minimal model of $\mathcal{K}$ s.t. $a^{\mathcal{I}} \in \mathsf{Subgoal}_2^{\mathcal{I}}$. Then, for every $d \in \mathsf{Node}^{\mathcal{I}}$, there exists a unique $t \in T$, that we denote $tile(d)$, such that $(d, a_t) \in \mathsf{tile}^{\mathcal{I}}$. Furthermore, for every $d, e \in \mathsf{Node}^{\mathcal{I}}$, the following properties hold:*

1. *if $(d, e) \in \mathsf{h}^{\mathcal{I}}$, then $(tile(d), tile(e)) \in H$;*
2. *if $(d, e) \in \mathsf{v}^{\mathcal{I}}$, then $(tile(d), tile(e)) \in V$;*
3. *if $pos(d) \in \mathsf{Pos} \setminus \{\mathsf{C}\}$, then $tile(d)$ has color b on the corresponding $pos(d)$-border(s).*

It now remains to address the main problem that is how to guarantee that, in minimal models satisfying Goal, the h-v-paths of nodes from a collapse into an actual grid. The idea is to force the satisfaction of a concept X somewhere along one of these h-v-paths. If this instance of X is placed on a node e where paths are forming the intended grid, that is e is both the h-v- and the v-h-successor of an element d, then it triggers a *flooding* concept Flood. The concept Flood then propagates along *all* h-v-paths and makes the concept X also satisfied everywhere, following the intuition already highlighted in Example 3. If a model does not feature a proper grid, then the flooding can be avoided by placing X somewhere the paths are not closing as a grid. Thus, a minimal model that does not feature a grid is not flooded. On the other hand, if a minimal model features a proper grid, then it is impossible to avoid the flooding and in particular, Flood holds at a. We can thus use the conjunction of Root and Flood as the final goal.

We now explain how to force the placement of X by 'guessing' an h-v-path of nodes from the root. Being along that path is represented by a concept PX, and we force the root a to satisfy this concept as soon as the previous subgoals are satisfied with the axiom:

$$\mathsf{Root} \sqcap \mathsf{Subgoal}_1 \sqcap \mathsf{Subgoal}_2 \sqsubseteq \mathsf{PX}$$

Now, if a node satisfies the concept PX, and depending on its own position, it either satisfies X or propagates the concept PX either horizontally or vertically. Three of the corner positions actually have no choice (*e.g.* the north-east corner

$$\exists\mathsf{pos}.p \sqsubseteq \mathsf{HGoodP} \quad \text{for } p \in \{\mathsf{SE},\mathsf{E},\mathsf{NE}\} \qquad \exists\mathsf{pos}.p \sqcap \exists\mathsf{h}.(\mathsf{GoodP} \sqcap \exists\mathsf{pos}.p') \sqsubseteq \mathsf{HGoodP} \quad \text{for } p,p' \in \mathsf{Pos} \text{ s.t. } p \rightsquigarrow_h p'$$
$$\exists\mathsf{pos}.p \sqsubseteq \mathsf{VGoodP} \quad \text{for } p \in \{\mathsf{NW},\mathsf{N},\mathsf{NE}\} \qquad \exists\mathsf{pos}.p \sqcap \exists\mathsf{v}.(\mathsf{GoodP} \sqcap \exists\mathsf{pos}.p') \sqsubseteq \mathsf{VGoodP} \quad \text{for } p,p' \in \mathsf{Pos} \text{ s.t. } p \rightsquigarrow_v p'$$
$$\mathsf{HGoodP} \sqcap \mathsf{VGoodP} \sqsubseteq \mathsf{GoodP} \qquad \mathsf{GoodP} \sqcap \mathsf{Root} \sqsubseteq \mathsf{Subgoal}_1$$

Figure 2: Additional rules to guarantee consistent positions.

cannot propagate any further thus must satisfy X):

$$\mathsf{PX} \sqcap \exists\mathsf{pos}.\mathsf{NE} \sqsubseteq \mathsf{X}$$
$$\mathsf{Node} \sqcap \mathsf{PX} \sqcap \exists\mathsf{pos}.\mathsf{NW} \sqsubseteq \exists\mathsf{h}.(\mathsf{Node} \sqcap \mathsf{PX})$$
$$\mathsf{Node} \sqcap \mathsf{PX} \sqcap \exists\mathsf{pos}.\mathsf{SE} \sqsubseteq \exists\mathsf{v}.(\mathsf{Node} \sqcap \mathsf{PX})$$

For the remaining positions, the possible options are represented by dedicated s-successors as follows:

$$\exists\mathsf{pos}.p \sqcap \mathsf{PX} \sqsubseteq \exists\mathsf{s}.\mathsf{isX} \quad \text{for } p \in \{\mathsf{C},\mathsf{N},\mathsf{E}\}$$
$$\exists\mathsf{pos}.p \sqcap \mathsf{PX} \sqsubseteq \exists\mathsf{s}.\mathsf{H} \quad \text{for } p \in \{\mathsf{C},\mathsf{N},\mathsf{S},\mathsf{W},\mathsf{SW}\}$$
$$\exists\mathsf{pos}.p \sqcap \mathsf{PX} \sqsubseteq \exists\mathsf{s}.\mathsf{V} \quad \text{for } p \in \{\mathsf{C},\mathsf{E},\mathsf{S},\mathsf{W},\mathsf{SW}\}.$$

The choice is then made via an extra s-successor that may or may not collapse with the possible options: for $p \in \mathsf{Pos} \setminus \{\mathsf{NE},\mathsf{NW},\mathsf{SE}\}$, consider the axioms:

$$\exists\mathsf{pos}.p \sqcap \mathsf{PX} \sqsubseteq \exists\mathsf{s}.\mathsf{Ch}$$
$$\exists\mathsf{s}.(\mathsf{Ch} \sqcap \mathsf{isX}) \sqsubseteq \mathsf{X}$$
$$\mathsf{Node} \sqcap \exists\mathsf{s}.(\mathsf{Ch} \sqcap \mathsf{H}) \sqsubseteq \exists\mathsf{h}.(\mathsf{Node} \sqcap \mathsf{PX})$$
$$\mathsf{Node} \sqcap \exists\mathsf{s}.(\mathsf{Ch} \sqcap \mathsf{V}) \sqsubseteq \exists\mathsf{v}.(\mathsf{Node} \sqcap \mathsf{PX})$$

Note that the two latter axioms could justify additional instances of roles h and v. However, due to Node $\sqcap$ PX being more specific than Node, Claim 1 still holds.

It is crucial that the (up to 3) s-successors above, carrying the different possibilities isX, H, V, do not collapse together. To prevent this, we use an error-detection mechanism that reports back to the root a, via the following axioms:

$$\exists\mathsf{s}.(\mathsf{isX} \sqcap \mathsf{H}) \sqsubseteq \mathsf{Err} \quad \exists\mathsf{s}.(\mathsf{isX} \sqcap \mathsf{V}) \sqsubseteq \mathsf{Err} \quad \exists\mathsf{s}.(\mathsf{H} \sqcap \mathsf{V}) \sqsubseteq \mathsf{Err}$$
$$\exists\mathsf{h}.\mathsf{Err} \sqsubseteq \mathsf{Err} \qquad \exists\mathsf{v}.\mathsf{Err} \sqsubseteq \mathsf{Err}.$$

If the error concept holds at a, we force the model to collapse in a way that cannot satisfy the final goal predicate. This is achieved by introducing an auxiliary element c that could act as an horizontal and vertical successor node for a, except that c misses the concept name Node. Consider the following assertions:

$$\mathsf{h}(a,c), \mathsf{v}(a,c), \mathsf{PX}(c), \mathsf{pos}(c,c), \mathsf{tile}(c,c), \mathsf{Any}(c), \mathsf{spy}(c,a).$$

Now the trick is to promote c to be a node whenever Err holds on a, which is achieved by the axiom $\exists\mathsf{spy}.\mathsf{Err} \sqsubseteq \mathsf{Node}$. Notice that the interpretation $\mathcal{I}_0$ obtained by interpreting every predicate as in the ABox, except for $\mathsf{Node}^{\mathcal{I}_0} := \{a, c\}$, is a model of $\mathcal{K}$. Therefore, in a minimal model $\mathcal{I}$, if the error predicate is to be seen along a h-v-path of nodes from a, then it triggers the concept Err on a, thus c is an instance of Node and therefore $\mathcal{I}_0 \subseteq \mathcal{I}$, thus $\mathcal{I} = \mathcal{I}_0$ by minimality of $\mathcal{I}$. We summarize this latter trick in the following claim:

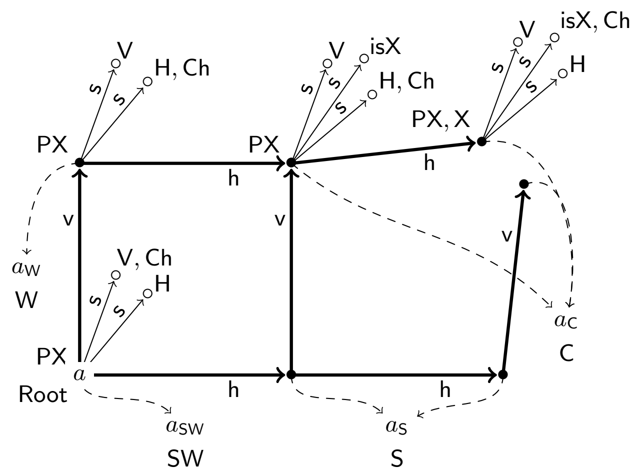

Figure 3: A part of the South-West corner of a model. Dashed arrows represent the role pos. Instances of the concept Node are a and the •-anonymous elements. This small portion fails to collapse as a proper 3×2 grid. Instances of PX can thus follow the defective v-h-h-path to place concept X in a way that avoids flooding.

Claim 4. *If $\mathcal{I}$ is a minimal model of $\mathcal{K}$, then* $(\mathsf{isX} \sqcap \mathsf{H})^{\mathcal{I}} = (\mathsf{isX} \sqcap \mathsf{V})^{\mathcal{I}} = (\mathsf{H} \sqcap \mathsf{V})^{\mathcal{I}} = \varnothing$.

We now trigger the flooding as previously described, with the axiom: $\exists\mathsf{h}.\exists\mathsf{v}.\mathsf{X} \sqcap \exists\mathsf{v}.\exists\mathsf{h}.\mathsf{X} \sqsubseteq \mathsf{Flood}$. Figure 3 illustrates how X can be placed to avoid triggering this latter rule if the model does not close as a grid. Now, propagating Flood back to a is easily performed by the axioms:

$$\exists\mathsf{h}.\mathsf{Flood} \sqsubseteq \mathsf{Flood} \qquad \exists\mathsf{v}.\mathsf{Flood} \sqsubseteq \mathsf{Flood}.$$

This propagates the flooding predicate to nodes along the h-v-path back to a, but not to *all* nodes along *all* h-v-paths. To achieve this latter part, recall that Point 3 of Claim 2 guarantees that, wherever we are on a path, there is always a further node with position NE, *i.e.* a node that is connected to a_{NE} by role pos. We can thus ensure that Flood propagates to every node by adding the following assertion and axiom:

$$\mathsf{aux}(a_{\mathsf{NE}}, a) \qquad \exists\mathsf{pos}.(\exists\mathsf{aux}.\mathsf{Flood}) \sqsubseteq \mathsf{Flood}.$$

As announced, we now require Flood to "flood" the model by making all the X-related concepts and choices satisfied:

$$\mathsf{Flood} \sqsubseteq \mathsf{X} \sqcap \mathsf{PX}$$
$$\exists\mathsf{pos}.p \sqcap \mathsf{Flood} \sqsubseteq \exists\mathsf{s}.(\mathsf{isX} \sqcap \mathsf{Ch}) \quad \text{for } p \in \{\mathsf{C},\mathsf{N},\mathsf{E}\}$$
$$\exists\mathsf{pos}.p \sqcap \mathsf{Flood} \sqsubseteq \exists\mathsf{s}.(\mathsf{H} \sqcap \mathsf{Ch}) \quad \text{for } p \in \{\mathsf{C},\mathsf{N},\mathsf{S},\mathsf{W},\mathsf{SW}\}$$
$$\exists\mathsf{pos}.p \sqcap \mathsf{Flood} \sqsubseteq \exists\mathsf{s}.(\mathsf{V} \sqcap \mathsf{Ch}) \quad \text{for } p \in \{\mathsf{C},\mathsf{E},\mathsf{S},\mathsf{W},\mathsf{SW}\}$$

We conclude the construction of the KB $\mathcal{K}$ by adding the axiom $\mathsf{Flood} \sqcap \mathsf{Root} \sqsubseteq \mathsf{Goal}$. Before proceeding to prove the reduction, it is useful to state one last property that summarizes some basic properties of minimal models of $\mathcal{K}$ and highlights the effect of the latter flooding axioms:

Claim 5. *Let $\mathcal{I}$ be a minimal model of $\mathcal{K}$. Then concept names from* $\{\mathsf{Root}\} \cup \mathsf{Pos} \cup T$*, as well as role names* spy *and* aux *are interpreted as specified in the ABox. Furthermore, if* $a^{\mathcal{I}} \in \mathsf{Goal}^{\mathcal{I}}$*, then* $\mathsf{Subgoal}_1^{\mathcal{I}} = \mathsf{Subgoal}_2^{\mathcal{I}} = \mathsf{Goal}^{\mathcal{I}} = \{a^{\mathcal{I}}\}$ *and* $\mathsf{Node}^{\mathcal{I}} = \mathsf{Flood}^{\mathcal{I}}$.

We can now prove that the reduction is correct.

Claim 6. $(T, b) \in$ RECTTILE *iff there exists a minimal model of $\mathcal{K}$ that satisfies* Goal.

Proof sketch. The $(\Rightarrow)$ direction is the easier one. From the tile assignment σ of a finite rectangle, one can define a corresponding model $\mathcal{I}$ that satisfies concept Goal and features an actual grid, with in particular $\mathsf{Node}^{\mathcal{I}} = \mathsf{Flood}^{\mathcal{I}}$. The minimality of $\mathcal{I}$, and notably the necessity of $\mathsf{Node}^{\mathcal{I}} = \mathsf{Flood}^{\mathcal{I}}$, is then established by arguing that the flooding cannot be avoided due to $\mathcal{I}$ encoding a proper grid. Thus, wherever X is placed, it always triggers the flooding.

The $(\Leftarrow)$ direction is more involved. From a minimal model $\mathcal{I}$ that satisfies Goal, we identify in $\mathcal{I}$ a valid tile assignment on a finite rectangle. To do so, we say that an element $d \in \mathsf{Node}^{\mathcal{I}}$ has coordinate (i, j) if there exists a h-v-path $d_0, r_1, d_1 \ldots, r_n, d_n$ of nodes from a with $d_n = d$ and such that there are exactly $i-1$ occurrences of h and $j-1$ of v in $r_1, \ldots, r_n$. Note that, by Claim 1, each $d \in \mathsf{Node}^{\mathcal{I}}$ has at least one such coordinate. We then check that if two elements in $\mathcal{I}$ share a coordinate, then they are equal, which is the key of our construction. This is argued by contradiction: assume there are two h-v-paths $d_0, r_1, d_1 \ldots, r_n, d_n$ and $e_0, s_1, e_1 \ldots, s_n, e_n$ of nodes from a, that both take the same numbers i of h-edges and j of v-edges, but $d_n \neq e_n$. We prove there is a $1 \leq k \leq n$ such that, placing the one necessary instance of X at d_k, we obtain a model $\mathcal{J}$ of $\mathcal{K}$ that avoids flooding, thus contradicting the minimality of $\mathcal{I}$. $\square$

4 Acyclicity to the Rescue

In the light of Theorem 1, is there any hope for minimal model reasoning in DLs? In our search for positive results, we look for inspiration in *pointwise circumscription*, where decidability results have been obtained for rather expressive fragments of $\mathcal{ALCIO}$ known to be undecidable in classical circumscription, notably including cases where roles are minimized. Pointwise circumscription coincides with standard circumscription for large classes of acyclic TBoxes, suggesting that terminological cycles may play a key role in the infeasibility of minimization. This is also supported by the heavy use of cyclic inclusions in our undecidability proof. We thus turn our attention to acyclicity notions, and find that the excursion is fruitful: minimal model reasoning becomes much more manageable for *acyclic TBoxes.*

4.1 Strong Acyclicity

Following (Di Stefano and Šimkus 2024), for an $\mathcal{ALCIO}$ concept C in negation normal form (NNF), we define the sets $Occ^+(C)$ and $Occ^-(C)$ of predicates that occur in C *positively* and *negatively*, respectively:

$$\begin{aligned}
Occ^+(A) = Occ^-(\neg A) &= \{A\} && \text{with } A \in \mathsf{N_C}\\
Occ^+(\neg A) = Occ^-(A) &= \varnothing && \text{with } A \in \mathsf{N_C}\\
Occ^+(\{a\}) = \{\{a\}\} \quad Occ^-(\{a\}) &= \varnothing && \text{with } a \in \mathsf{N_I}\\
Occ^\pm(C \circ D) &= Occ^\pm(C) \cup Occ^\pm(D) && \circ \in \{\sqcup, \sqcap\}\\
Occ^+(\exists r.C) &= \{r\} \cup Occ^+(C)\\
Occ^-(\exists r.C) &= Occ^-(C)\\
Occ^+(\forall r.C) &= Occ^+(C)\\
Occ^-(\forall r.C) &= \{r\} \cup Occ^-(C)
\end{aligned}$$

We let $Occ^\pm(C \sqsubseteq D) = Occ^\pm(NNF(\neg C \sqcup D))$ for a concept inclusion $C \sqsubseteq D$ and $\pm \in \{+, -\}$. The *dependency graph* $DG(\mathcal{T})$ of an $\mathcal{ALCIO}$ TBox $\mathcal{T}$ has as nodes all the concept names and role names and all concepts of the form $\{a\}$ or $\top$ that appear in $\mathcal{T}$, and there is an edge from P_1 to P_2 if $P_1 \in Occ^-(\alpha)$ and $P_2 \in Occ^+(\alpha)$ for some $\alpha \in \mathcal{T}$. We 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 (Baader and Nutt 2003) which is satisfied, for example, by the well-known medical terminology SNOMEDCT. Under the classical semantics, terminologies often rely on concept-equivalences and are therefore cyclic TBoxes. If we take the definitions in a terminology $A \doteq C$ as inclusions $C \sqsubseteq A$, we may regain acyclicity (and may enjoy lower complexity). Example 4 illustrates that, under the minimal model semantics, keeping only one of the two directions might be innocuous since the other inclusion is enforced by predicate minimization.

Example 4. *A patient has been diagnosed with pneumoconiosis, which can be caused by various organic dust types; some types are very serious.* Baritosis *is caused by barium dust, as stated in the following Snomed CT definition:*[1]

$$\mathsf{Pneumoconiosis} \sqcap \exists caus\_agent.\mathsf{Barium\_Dust} \sqsubseteq \mathsf{Baritosis}$$

Under the classical semantics, there are models where the patient is diagnosed with baritosis due to a causative agent, barium dust, that is not justified in a real finding but simply made true to cause the baritosis diagnosis. In the minimal models, in contrast, baritosis can only be diagnosed on the basis of justified clinical findings.

Standard circumscription has been studied for acyclic terminologies (Bonatti, Lutz, and Wolter 2009), but unlike in our setting, acyclicity does not help reducing the complexity. Reasoning about general $\mathcal{ALCIO}$ KBs can be reduced to acyclic ones, but the reduction uses *varying* predicates.

Despite their restrictions, acyclic TBoxes are quite expressive, and strongly acyclic $\mathcal{EL}$ can force minimal models satisfying a concept of interest to have an exponential size, as illustrated by Example 5. We later show that this exponential size is also sufficient, in the sense that every concept satisfiable in a minimal model also is in one with exponential size, and this even for the relaxed notion of weak acyclicity.

[1] Based on https://pmc.ncbi.nlm.nih.gov/articles/PMC4422531/

Example 5. *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 enforce all leaves to be different objects in minimal models that satisfy a concept of interest. 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 two assertions $\mathsf{Left}(o)$ and $\mathsf{Right}(o')$ and the following axioms for $0 \leq j < n$:*

$$\mathsf{L}'_j \sqsubseteq \exists \mathsf{pick}.\top \qquad \mathsf{L}'_j \sqcap \exists \mathsf{pick}.\mathsf{Left} \sqsubseteq \exists \mathsf{l}'_j.\mathsf{L}'_{j+1,l}$$
$$\mathsf{L}'_{j+1,l} \sqcap \mathsf{L}'_{j+1,r} \sqsubseteq \mathsf{L}'_{j+1} \qquad \mathsf{L}'_j \sqcap \exists \mathsf{pick}.\mathsf{Right} \sqsubseteq \exists \mathsf{r}'_j.\mathsf{L}'_{j+1,r}$$

A minimal model of this strongly acyclic $\mathcal{EL}$ KB can only satisfy the concept L'_n (representing the root of the second tree) if its interpretation of the first tree produces at least 2^n instances of L_n, i.e. of L'_0.

Before moving to the complexity results, we provide another illustrative example: in strongly acyclic $\mathcal{EL}$ we can assign truth values to objects and compare these assignments. This trick will be useful in the hardness proofs below.

Example 6. *Consider the ABox with assertions $\mathsf{N}_1(a)$, $\mathsf{N}_2(b)$, $\mathsf{T}(v_1)$, $\mathsf{F}(v_2)$, and the strongly acyclic $\mathcal{EL}$ TBox with the following axioms, for $i \in \{1, 2\}$ and $C \in \{\mathsf{T}, \mathsf{F}\}$:*

$$\mathsf{N}_i \sqsubseteq \exists \mathsf{val}.\mathsf{TV} \qquad \mathsf{N}_1 \sqsubseteq \exists \mathsf{read}.\top$$
$$\exists \mathsf{val}.C \sqsubseteq C' \qquad C' \sqcap \exists \mathsf{read}.(\mathsf{N}_2 \sqcap C') \sqsubseteq \mathsf{Goal}$$

The nodes a and b each pick a truth value via role val*, which is copied to the node as T' (resp. F')* only if *it is T (resp. F). The node a then reads the value at some object via role* read*, and the* Goal *concept is satisfied exactly when a reads the value at b (the only instance of N_2 in a minimal model is b) and they both picked the same truth value.*

Strongly Acyclic $\mathcal{EL}$ and Pointwise Circumscription. To show the decidability of strongly acyclic $\mathcal{ELIO}_\bot$ we rely on results on *pointwise circumscription* (Di Stefano, Ortiz, and Simkus 2023), a local approximation of standard circumscription where minimization is allowed only locally, at one domain element. In our definition of minimal models, predicates are minimized *globally*, across the entire interpretation. In pointwise circumscription, we refine the relation $\subseteq$ by only considering *pointwise comparable* interpretations.

Definition 2 (Pointwise Comparison)**.** *Given interpretations $\mathcal{I}$ and $\mathcal{J}$, we write $\mathcal{I} \sim^\bullet \mathcal{J}$ if there exists $e \in \Delta^\mathcal{I}$ such that:*

(i) for all $A \in \mathsf{N_C}$, $A^\mathcal{I} \cap \Delta = A^\mathcal{J} \cap \Delta$, and
(ii) for all $r \in \mathsf{N_R}$, $r^\mathcal{I} \cap (\Delta \times \Delta) = r^\mathcal{J} \cap (\Delta \times \Delta)$,

where $\Delta = \Delta^\mathcal{I} \setminus \{e\}$.

Definition 3 (Pointwise Minimal Model)**.** *Given a KB $\mathcal{K}$, a model $\mathcal{I}$ is pointwise minimal if there exists no $\mathcal{J} \subsetneq \mathcal{I}$ such that $\mathcal{J} \models \mathcal{K}$ and $\mathcal{J} \sim^\bullet \mathcal{I}$.*

To emphasize the difference, we sometimes refer to minimal models in the sense of this paper as *globally* minimal models. In general, the set of pointwise minimal models does not coincide with the set of globally minimal models.

Example 7. *Let $\mathcal{K} := (\{\exists \mathsf{r}.\mathsf{A} \sqsubseteq \mathsf{A}\}, \{\mathsf{r}(a, b), \mathsf{r}(b, a)\})$. The interpretation $\mathcal{I}$ such that $\Delta^\mathcal{I} = \{d, e\}$, $a^\mathcal{I} = d$ and $b^\mathcal{I} = e$, $\mathsf{A}^\mathcal{I} = \{d, e\}$ and $\mathsf{r}^\mathcal{I} = \{(d, e), (e, d)\}$ is a pointwise minimal model of our $\mathcal{EL}$ KB $\mathcal{K}$ that is not globally minimal.*

Pointwise circumscription has better computational properties than standard circumscription. The modal depth of a KB $\mathcal{K}$, denoted with $\mathrm{md}(\mathcal{K})$, is defined as the maximal number of nested quantifiers occurring in $\mathcal{K}$. When all roles are minimized, concept satisfiability is complete for NExp in the fragment $\mathcal{ALCIO}^{d\leq 1}$ of $\mathcal{ALCIO}$ with modal depth one (Di Stefano, Ortiz, and Simkus 2023). In contrast, the problem is undecidable for full $\mathcal{ALCIO}$.

Standard normalization techniques (Baader et al. 2017) do not preserve minimal models of $\mathcal{ALCIO}$ KBs in general, but they do for $\mathcal{ELIO}_\bot$ KBs.

Proposition 1. *Any KB $\mathcal{K}$ in $\mathcal{ELIO}_\bot$ can be transformed in polynomial time into an KB $\mathcal{K}'$ in $\mathcal{ELIO}_\bot$ with $\mathrm{md}(\mathcal{K}') \leq 1$ such that $\mathcal{K}'$ is a conservative extension of $\mathcal{K}$ under the minimal model semantics; moreover $\mathcal{K}'$ preserves strong acyclicity and weak acyclicity (the latter is defined in Section 4.2).*

For strongly acyclic KBs in $\mathcal{ALCIO}^{d\leq 1}$, minimal models and pointwise minimal models coincide (Di Stefano and Šimkus 2024). We thus inherit the following result.

Theorem 2. MINMODSAT *in strongly acyclic $\mathcal{ALCIO}^{d\leq 1}$ is in* NExp.

The complexity result of Theorem 2 applies to strongly acyclic $\mathcal{ELIO}_\bot$ without restrictions on the modal depth.

Theorem 3. MINMODSAT *in strongly acyclic $\mathcal{ELIO}_\bot$ is in* NExp.

Proof. Let $\mathcal{ELIO}_\bot^{d\leq 1}$ be the fragment of $\mathcal{ELIO}_\bot$ where concept expressions have *modal depth* at most one. Pointwise minimal satisfiability in $\mathcal{ELIO}_\bot^{d\leq 1}$ is in NExp. In strongly acyclic $\mathcal{ELIO}_\bot$, the set of pointwise minimal models coincides with the set of minimal models. The claim then follows from Proposition 1. ☐

Theorem 4. MINMODSAT *in strongly acyclic $\mathcal{EL}$ is* NExp*-hard.*

Proof sketch. We provide a reduction from the *torus tiling problem* (Tobies 1999) to MINMODSAT. Using the construction in Example 5, we construct a $\mathcal{K}$ and a goal concept Goal such that the satisfaction of Goal in a minimal model $\mathcal{I}$ ensures that: (1) in $\mathcal{I}$ we can embed a tree of depth $2n$, where each leaf encodes (in binary) a pair of coordinates (x, y), with $0 \leq x, y \leq 2^n - 1$; (2) a torus is embedded in the leaves of the tree in a way such that the horizontal and vertical successors respect the tiling conditions. To achieve the latter desiderata, we construct a subgoal concept G_{2n} ensuring that each leaf encoding the pair (x, y) has as horizontal successor the pair $(x+1, y)$ and as vertical successor the pair $(x, y+1)$. To check that G_{2n} is satisfied at all the leaves, we propagate a concept $\mathsf{LeafGrid}$ back to the root, using the following axioms:

$$\exists \mathsf{l}_i.\mathsf{G}_{i+1} \sqcap \exists \mathsf{r}_i.\mathsf{G}_{i+1} \sqsubseteq \mathsf{G}_i \quad \text{with } 0 \leq i < 2n$$
$$\mathsf{G}_0 \sqsubseteq \mathsf{LeafGrid}$$

where the roles l_i and r_i are as in Example 5. In a minimal model, $\mathsf{LeafGrid}$ is satisfied at the root of the tree if G_{2n} is satisfied at all the leaves. □

4.2 Weak Acyclicity

We now define *weak acyclicity*, which is an important notion for TGDs in the database literature. We refine the above notion by annotating some edges in $DG(\mathcal{T})$ as $\star$-edges. For a concept C in NNF, we define:

$$\begin{aligned}
Occ^{+}_{\exists}(A) &= Occ^{+}_{\exists}(\neg A) = \varnothing && \text{with } A \in \mathsf{N_C}\\
Occ^{+}_{\exists}(\{o\}) &= \varnothing && \text{with } a \in \mathsf{N_I}\\
Occ^{+}_{\exists}(C \circ D) &= Occ^{+}_{\exists}(C) \cup Occ^{+}_{\exists}(D) && \circ \in \{\sqcup, \sqcap\}\\
Occ^{+}_{\exists}(\exists r.C) &= \{r\} \cup Occ^{+}(C)\\
Occ^{+}_{\exists}(\forall r.C) &= Occ^{+}_{\exists}(C)
\end{aligned}$$

We let $Occ^{+}_{\exists}(C \sqsubseteq D) = Occ^{+}_{\exists}(NNF(\neg C \sqcup D))$ for a concept inclusion $C \sqsubseteq D$. In $DG(\mathcal{T})$, there is a $\star$-edge from P_1 to P_2 if for some $\alpha \in \mathcal{T}$, $P_1 \in Occ^{-}(\alpha)$ and $P_2 \in Occ^{+}_{\exists}(\alpha)$. Notice that for all C, we have $Occ^{+}_{\exists}(C) \subseteq Occ^{+}(C)$, and thus all $\star$-edges are also basic edges. We call a TBox $\mathcal{T}$ *weakly acyclic* if there is no cycle in $DG(\mathcal{T})$ that goes through an $\star$-edge and no node is reachable from $\top$ in $DG(\mathcal{T})$. Clearly, every strongly acyclic TBox is also weakly acyclic. The $\mathcal{EL}$ TBox in Example 7 is weakly acyclic but not strongly acyclic.

Fortunately, even for weakly acyclic $\mathcal{ELIO}_\perp$, we can always find a model whose size is at most single exponential.

Lemma 1. *Weakly acyclic $\mathcal{ELIO}_\perp$ has the small model property: if a concept C is satisfied in some minimal model of $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, then it is satisfied in a minimal model $\mathcal{J}$ whose domain has size bounded by $|\mathsf{N_I}(\mathcal{K})|(|\mathcal{T}|2^{|\mathcal{T}|})^{|\mathcal{T}|}$.*

Proof sketch. Let $\mathcal{I}$ be a minimal model of $\mathcal{K}$ satisfying a concept C. It suffices to prove that its active domain *i.e.* the subset of elements from $\Delta^{\mathcal{I}}$ that occur in the interpretation of at least one concept or role name, has the claimed size. Since no node is reachable from $\top$ in $DG(\mathcal{T})$, every single fact in $\mathcal{I}$ somehow stems from the individuals occurring in $\mathcal{K}$. We thus start from those and track which successors they might require based on their types, *i.e.* the combinations of concepts they satisfy in $\mathcal{I}$. For example, an element a with type $\{A\}$ requires an r-successor with type containing B if $\mathcal{T}$ has an axiom $A \sqsubseteq \exists r.B$; since $\mathcal{I}$ is a model, there exists such a successor e. Now, instead of directly iterating the above by looking at the type of e in $\mathcal{I}$ (say, $\{A, B\}$), we first restrict this type to the concepts that are reachable from A via *at least* one $\star$-edge of $DG(\mathcal{T})$ (for element e, we thus restrict its type $\{A, B\}$ to $\{B\}$). Indeed, no such concept can further require A: it would form a cycle in $DG(\mathcal{T})$ going through the said $\star$-edge, contradicting $\mathcal{T}$ being weakly acyclic. Therefore, the restricted types we successively consider become empty after at most $|DG(\mathcal{T})|$ iterations. In particular, the number of reached elements is bounded as desired; and the restriction of $\mathcal{I}$ to these elements is a model, which must be the active domain of $\mathcal{I}$ as $\mathcal{I}$ is minimal. □

This lemma is our key to deriving tight complexity bounds for the weakly acyclic setting.

Theorem 5. MINMODSAT *in weakly acyclic $\mathcal{ELIO}_\perp$ is $\mathsf{NExp}^{\mathsf{NP}}$-complete. The lower bound holds already for $\mathcal{EL}$.*

Proof sketch. The upper bound immediately follows from Lemma 1. Indeed, we can use a naive procedure that "guesses" an exponentially large candidate model $\mathcal{I}$ of the input KB, and checks for the non-existence of a model $\mathcal{J} \subsetneq \mathcal{I}$ using an NP oracle.

For the lower bound, we provide a reduction from (the complement of) *succinct* CERT3COL (Eiter, Gottlob, and Mannila 1997) to MINMODSAT. Our reduction combines the ideas behind Example 3, using flooding, and those illustrated in Examples 5 and 6 that allow us to succinctly represent the exponentially large graph. We construct a weakly acyclic KB $\mathcal{K}$ in $\mathcal{EL}$ and define some *subgoal* concepts which are needed for the goal concept to be satisfied, and use them as in Example 5 to ensure that in every minimal model we can find the following trees:

- A tree T_C per each color $C \in \{R, G, B\}$ of depth n, where each leaf corresponds to a vertex of the input graph with a color assignment. Using a minimality argument, we ensure that such trees have disjoint sets of leaves.
- A tree T_G of depth n where each leaf corresponds to a node of the input graph. We craft a subgoal concept Col that must be satisfied at each leaf of T_G. The subgoal Col ensures that (a) each leaf in T_G is connect to the leaf in T_C corresponding to the same node in the graph, for each $C \in \{R, G, B\}$; (b) at least one of the leaves is marked as the chosen color.
- A tree T_V of depth $2n$ where each leaf corresponds to a variable $v_{i,j}$. By crafting a dedicated subgoal concept, that must be satisfied at each leaf of T_V, we ensure that each variable has a unique truth assignment.
- A large tree T_F, of depth $6n + 2$, where each leaf encodes a tuple $(u, v, x, y, \sigma^1, \sigma^2)$ where: u, v are vertexes in the input graph G, x, y are variables and $\sigma^1, \sigma^2 \in \{0, 1\}$ (polarities of the variables). Each $(u, v, x, y, \sigma^1, \sigma^2)$ represents a possible edge in the input graph G. Simulating the computation of a family of circuits (following (Bonatti, Lutz, and Wolter 2009)), we mark the leaves of T_F corresponding to *real* edges in the graph G.

We construct two subgoal concepts that are satisfied if each leaf encoding $(u, v, x, y, \sigma^1, \sigma^2)$ connects via dedicated roles to the leaves corresponding to u, v (in T_G) and x, y (in T_V), as in Example 6. Using these role connections, the leaves in T_F import the color and truth assignments. By evaluating the clause encoded in each $(x, y, \sigma^1, \sigma^2)$, we mark the leaves in T_F corresponding to *true* edges in G. If a bad color assignment is detected, the structure gets *flooded* similarly to Example 3. In particular, all the leaves of T_C, for each $C \in \{R, G, B\}$ are marked as chosen. A concept Flood is then propagated at all the leaves of T_F. In particular, the satisfaction of Flood forces the leaves to be assigned to all colors. Instead of checking the satisfaction of each subgoal Goal at the (exponentially many) leaves, we transfer the check to the root of the corresponding tree as follows:

$$\exists \mathsf{l}_j.\mathsf{Goal} \sqcap \exists \mathsf{r}_j.\mathsf{Goal} \sqsubseteq \mathsf{Goal} \text{ for all } 0 \leq j \leq \text{tree depth}$$

In a minimal model, the subgoal Goal is satisfied at the root of the tree if and only if Goal is true at each leaf. We construct a final goal concept that is satisfied in a minimal model if all the roots of the different trees satisfy their respective goals. We emphasize that the propagation of the Flood concept requires axioms that are not strongly acyclic. □

Data Complexity We also look at the data complexity of MinModSat under acyclicity restrictions.

Theorem 6. MinModSat *for weakly acyclic $\mathcal{ELIO}_\bot$ is Σ_2^P-complete in data complexity. The lower bound applies already to $\mathcal{EL}$.*

Proof sketch. For the upper bound we use Lemma 1. Assuming that the TBox is fixed, if there exists a minimal model $\mathcal{I}$ of $(\mathcal{T}, \mathcal{A})$ that satisfies a concept C of interest, then there exists such an interpretation $\mathcal{J}$ whose domain is bounded by $c \times |\mathsf{N_I}(\mathcal{K})|$, where c is a constant that only depends on $\mathcal{T}$. In other words, the size of $\mathcal{J}$ that witnesses C is polynomial in the size of $\mathcal{A}$. Note that, given a candidate $\mathcal{J}$ as above, we can use NP oracle to check whether $\mathcal{J}$ is (non-)minimal. This yields the Σ_2^P upper bound.

For the lower-bound, we can mainly use the reduction that was described in Example 3. Specifically, we need to simulate the inclusions $V \sqsubseteq T \sqcup F$ and $N \sqsubseteq C_1 \sqcup C_2 \sqcup C_3$, which are not allowed in $\mathcal{EL}$. Instead, we can use inclusions of the form $A \sqsubseteq \exists r.B$, which contain a "hidden" disjunction via existential quantification. □

5 Minimal Models in Related Formalisms

We make a very brief excursion into the DL-Lite family, and briefly discuss our results in the setting of databases with *tuple-generating dependencies (TGDs)*. We also look at the impact of the UNA on minimal model reasoning in $\mathcal{EL}$.

5.1 DL-Lite

We did not study DL-Lite in this paper, and the feasibility of MinModSat in this family of DLs remains an intriguing question for future work. We only present one interesting result that hints that the problem will not be easy. In very stark contrast to the previously known NL-membership for MinModSat in DL-Lite$_{\text{core}}$ (Bonatti et al. 2023), already in DL-Lite$_{\text{horn}}$ we have ExpSpace-hardness.

Theorem 7. MinModSat *in* DL-Lite$_{\text{horn}}$ *is* ExpSpace-*hard.*

We believe that this bound, proved by reducing the acceptance problem of a Turing machine with exponential space, is likely to be tight, but leave the question for future work.

5.2 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., (Fagin et al. 2005; Calì, Gottlob, and Pieris 2012). 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 (Alviano, Morak, and Pieris 2017), 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 $\Sigma \cup \{\neg g(\vec{t}) \to \bot\}$, 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 (Alviano, Morak, and Pieris 2017), that relies on predicates of arity > 2, for weakly acyclic TGDs with stable negation.

5.3 $\mathcal{EL}$ without the UNA

Finally, we make an interesting observation about the role of the UNA in the presented results. Our hardness proofs all rely on the UNA, and use at least two individuals that must be interpreted as different objects in the domain. This is no coincidence: if we drop the UNA, MinModSat in $\mathcal{ELIO}$ and $\mathcal{EL}$ is not only decidable, it is even tractable.

Theorem 8. MinModSat *in $\mathcal{ELIO}$ is* P*-complete; the lower bound holds already for $\mathcal{EL}$.*

Even without the UNA, some concepts may not be satisfiable in a minimal model, but an $\mathcal{ELIO}$ KB now has a 'representative' minimal model with just one element. This representative model can be computed in polynomial time via a fixpoint computation akin to building a minimal model of a propositional definite Horn logic program. The lower bound also follows easily from the latter setting.

6 Conclusion

We have explored the challenges of reasoning with minimal models in description logics, and shown that enforcing minimality across all predicates leads to undecidability even in the lightweight $\mathcal{EL}$. This directly implies that minimal model reasoning in very restricted fragments of guarded TGDs with predicate arities ≤ 2 is also undecidable. Strong and weak acyclicity conditions allowed us to regain decidability and establish tight bounds on combined and data complexity in $\mathcal{EL}$, $\mathcal{ELIO}$, and even a fragment of $\mathcal{ALCIO}$. Some of these bounds are inherited for the recently studied setting of pointwise circumscription, providing further evidence that local, pointwise minimization is about the best we can do if we are interested in minimal models in DLs. It remains to be explored whether acyclicity conditions and pointwise minimization might also be useful in the richer setting of TGDs. One of the most intriguing avenues left open for further investigation is DL-Lite: we know that minimization is almost for free in DL-Lite$_{\text{core}}$, but makes reasoning ExpSpace-hard for its extension DL-Lite$_{\text{horn}}$. 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.

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

References

Alviano, M.; Morak, M.; and Pieris, A. 2017. Stable model semantics for tuple-generating dependencies revisited. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, 377–388.

Angiulli, F.; Ben-Eliyahu-Zohary, R.; Fassetti, F.; and Palopoli, L. 2014. On the tractability of minimal model computation for some CNF theories. *Artificial Intelligence* 210:56–77.

Baader, F., and Nutt, W. 2003. Basic description logics. In *Description Logic Handbook*. Cambridge University Press. 43–95.

Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.

Ben-Eliyahu, R., and Dechter, R. 1996. On computing minimal models. *Annals of Mathematics and Artificial Intelligence* 18(1):3–27.

Ben-Eliyahu-Zohary, R., and Palopoli, L. 1997. Reasoning with minimal models: Efficient algorithms and applications. *Artificial Intelligence* 96(2):421–449.

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

Bonatti, P. A.; Lutz, C.; and Wolter, F. 2009. The complexity of circumscription in DLs. *Journal of Artificial Intelligence Research* 35:717–773.

Calì, A.; Gottlob, G.; and Pieris, A. 2012. Towards more expressive ontology languages: The query answering problem. *Artificial Intelligence* 193:87–128.

Di Stefano, F., and Šimkus, M. 2024. 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)*, 306–316.

Di Stefano, F.; Manière, Q.; Ortiz, M.; and Šimkus, M. 2025. Minimal model reasoning in description logics: Don't try this at home! arxiv.org/abs/2508.05350.

Di Stefano, F.; Ortiz, M.; and Simkus, M. 2023. Description logics with pointwise circumscription. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 3167–3175.

Eiter, T., and Gottlob, G. 1995. On the computational cost of disjunctive logic programming: Propositional case. *Annals of Mathematics and Artificial Intelligence* 15(3-4):289–323.

Eiter, T.; Gottlob, G.; and Mannila, H. 1997. Disjunctive datalog. *ACM Trans. Database Syst.* 22(3):364–418.

Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336(1):89–124.

Gelfond, M., and Lifschitz, V. 1988. The stable model semantics for logic programming. In *Proceedings of the 5th International Logic Programming Conference and Symposium (ICLP/SLP)*, 1070–1080.

Grau, B. C.; Horrocks, I.; Krötzsch, M.; Kupke, C.; Magka, D.; Motik, B.; and Wang, Z. 2013. Acyclicity notions for existential rules and their application to query answering in ontologies. *Journal of Artificial Intelligence Research* 47:741–808.

Lackner, M., and Pfandler, A. 2012. Fixed-parameter algorithms for finding minimal models. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 85–95.

Lutz, C.; Manière, Q.; and Nolte, R. 2023. Querying Circumscribed Description Logic Knowledge Bases. In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 482–491.

McCarthy, J. 1980. Circumscription: a form of non-monotonic reasoning. *Artificial intelligence* 13(1-2):27–39.

Pfandler, A.; Pichler, R.; and Woltran, S. 2014. The complexity of handling minimal solutions in logic-based abduction1. *Journal of Logic and Computation* 25(3):805–825.

Reiter, R. 1980. A logic for default reasoning. *Artificial intelligence* 13(1-2):81–132.

Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32(1):57–95.

Stewart, I. A. 1991. Complete problems involving boolean labelled structures and projection transactions. *Journal of Logic and Computation* 1(6):861–882.

Tobies, S. 1999. A NExpTime-Complete Description Logic Strictly Contained in C^2. In *Proceedings of the Annual Conference for Computer Science Logic (CSL)*, 292–306.

Yang, J. 2014. Rectangular tileability and complementary tileability are undecidable. *European Journal of Combinatorics* 41:20–34.