Decidability of Querying First-Order Theories via Countermodels of Finite Width (extended abstract of journal article published in LMCS 2025)

Thomas Feller, Tim S. Lyon, Piotr Ostropolski-Nalewaja, Sebastian Rudolph
Computational Logic Group, TU Dresden, Germany

original paper reference: lmcs:11733 – Logical Methods in Computer Science, April 22, 2025, Volume 21, Issue 2, https://doi.org/10.46298/lmcs-21(2:7)2025

The notion of logical entailment is at the core of logic in computer science. Given that entailment in many natural logical formalisms is undecidable, much research has focused on identifying decidable cases or even finding generic principles underlying the divide between decidability and undecidability. The article summarized here contributes to this endeavor. Formally, we consider the following setting: Assuming two classes $\mathfrak{F}$ and $\mathfrak{Q}$ of logical sentences, the decision problem of interest is specified as follows:

Problem: $\textsc{Entailment}(\mathfrak{F}, \mathfrak{Q})$
Input: $\Phi \in \mathfrak{F}$, $\Psi \in \mathfrak{Q}$.
Output: YES, if $\Phi \models \Psi$, NO otherwise.

We refer to $\mathfrak{F}$ as a *specification language*, to $\mathfrak{Q}$ as a *query language*, and sometimes call the above entailment problem "query entailment" or "querying", as is common in database theory and knowledge representation. We assume $\mathfrak{F}$ and $\mathfrak{Q}$ both rely on a common model theory based upon relational structures of arbitrary cardinality, and focus on the setting where $\mathfrak{F}$ is a fragment of first-order logic (FO) and $\mathfrak{Q}$ is a fragment of universal second-order logic (∀SO).

The described problem in its unconstrained form, $\textsc{Entailment}(\text{FO}, \forall\text{SO})$, is merely semidecidable, that is, there exists an algorithm that successfully terminates in case $\Phi \models \Psi$. Thus, in order to arrive at a decidable setting, one needs to further constrain $\mathfrak{F}$ and $\mathfrak{Q}$ in a way that ensures co-semidecidability, i.e. the existence of a procedure successfully terminating whenever $\Phi \not\models \Psi$. We propose to achieve this by exploiting the existence of *countermodels* (that is, models of Φ that are not models of Ψ and hence witness the non-entailment) which are of a particularly simple shape, as indicated by having finite "width" according to appropriate width measures. As an easy example, consider the folklore observation that the satisfiability problem of a FO fragment $\mathfrak{F}$ is decidable whenever $\mathfrak{F}$ exhibits the *finite-model property*, meaning that every satisfiable sentence $\Phi \in \mathfrak{F}$ has a finite model. In this particular case, the used width measure would just be the domain size.

The article summarized here investigates under which conditions decidability of $\textsc{Entailment}(\mathfrak{F}, \mathfrak{Q})$ follows from its "*finite-width-countermodel property*" for some kind of width. To emphasize the relevance for KR, we will pair the presentation of the generic technical results with specific insights they provide for the field of description logics (DLs).

Width-Based Decidability. A *width measure* is a function w mapping every interpretation to a value from $\mathbb{N} \cup \{\infty\}$. An interpretation $\mathcal{I}$ is called w-finite if $\mathsf{w}(\mathcal{I}) \in \mathbb{N}$. Given a logical language $\mathfrak{L}$, a width measure w is called $\mathfrak{L}$-*friendly* if there exists an algorithm that, given a number $n \in \mathbb{N}$, enumerates all $\mathfrak{L}$-sentences that have a model $\mathcal{I}$ with $\mathsf{w}(\mathcal{I}) \leq n$. Particularly interesting instances are the GSO-friendly treewidth and MSO-friendly partitionwidth (see later; GSO/MSO stands for guarded/monadic second-order logic). A specification language $\mathfrak{F}$ is called w-*controllable* for a query language $\mathfrak{Q}$ if for every $\Psi \in \mathfrak{Q}$ and $\Phi \in \mathfrak{F}$ with $\Phi \not\models \Psi$, there exists a w-finite interpretation $\mathcal{I}$ with $\mathcal{I} \models \Phi$ and $\mathcal{I} \not\models \Psi$.

Theorem 1. *Let* w *be a width measure,* $\mathfrak{Q}$ *a* ∀*SO fragment, and* $\mathfrak{F}$ *an FO fragment that is* w*-controllable for* $\mathfrak{Q}$*. If the fragment* $\{\Phi \wedge \neg\Psi \mid \Phi \in \mathfrak{F}, \Psi \in \mathfrak{Q}\}$ *can be effectively expressed in a language* $\mathfrak{L}$ *for which* w *is* $\mathfrak{L}$*-friendly, then* $\textsc{Entailment}(\mathfrak{F}, \mathfrak{Q})$ *is decidable.*

Example 2. Rudolph and Glimm's proof of the previously long-standing open problem of decidability entailment of unions of conjunctive queries (UCQs) in the description logic $\mathcal{ALCHOIQ}b$ is an application of Theorem 1 with $\mathfrak{F} = \mathcal{ALCHOIQ}b$ and $\mathfrak{Q} = \text{UCQ}$, where $\mathfrak{L} = \text{FO}$ and treewidth is the used width notion w. The most intricate part of this result amounts to showing that $\mathcal{ALCHOIQ}b$ is treewidth-controllable for UCQ, via an elaborate countermodel construction, producing so-called *quasi-forest models*, which are known to be treewidth-finite.

Homomorphism-Closed Queries. A sentence Ψ is called *homomorphism-closed* if any interpretation $\mathcal{J}$ receiving a homomorphism from any model $\mathcal{I} \models \Psi$ is a model of Ψ, too.

Given a sentence Φ, a (sub)set $\mathscr{I}$ of its models is called a *finitely universal model set* of Φ, if for every model $\mathcal{J} \models \Phi$ there is some $\mathcal{I} \in \mathscr{I}$ such that every finite substructure of $\mathcal{I}$ maps homomorphically into $\mathcal{J}$.

Generalizing folklore knowledge about universal models, we show that in such a setting, $\Phi \not\models \Psi$ holds exactly if there exists a model $\mathcal{I} \in \mathscr{I}$ with $\mathcal{I} \not\models \Psi$. This insight ensures that for homomorphism-closed queries, the search for countermodels can be confined to any finitely universal model set of Φ. That means, the existence of a finitely universal model set that consists only of structurally well-behaved models warrants controllability. Together with Theorem 1, this allows us to establish the following theorem.

Theorem 3. *Let* w *be a width measure,* $\mathfrak{Q}$ *a* $\forall$*SO fragment of homomorphism-closed sentences, and* $\mathfrak{F}$ *an FO fragment whose every sentence has a finitely universal model set of only* w*-finite instances. If the fragment* $\{\Phi \wedge \neg\Psi \mid \Phi \in \mathfrak{F}, \Psi \in \mathfrak{Q}\}$ *can be effectively expressed in a language* $\mathfrak{L}$ *for which* w *is* $\mathfrak{L}$*-friendly, then* ENTAILMENT$(\mathfrak{F}, \mathfrak{Q})$ *is decidable.*

Example 4. For a variety of very expressive DLs – notably $\mathcal{ALCHOI}b$, $\mathcal{ALCHIQ}b$, $\mathcal{ALCHOQ}b$, and their sublogics – specific techniques commonly known as *unravelling* can transform any model $\mathcal{I}$ into a treewidth-finite model $\mathcal{J}$ that homomorphically maps into $\mathcal{I}$. Thus, due to GSO being treewidth-friendly, Theorem 3 ensures in one sweep decidability of entailment of all effectively GSO-expressible homomorphism-closed $\forall$SO queries for all these DLs.[1]

While for simple instances thereof like CQs or queries with regular path expressions, decidability of ENTAILMENT$(\mathfrak{F}, \mathfrak{Q})$ has been known for all these instances of $\mathfrak{F}$, our result not only provides a generic common ground for these prior results, but also generalizes them to much more expressive query languages like frontier-guarded disjunctive datalog and nested monadically defined queries.

Partitionwidth. We investigate a variant of Blumensath's *partitionwidth* – whose technical definiton we omit here for space reasons – and find it particularly useful for our purposes. Partitionwidth-finiteness generalizes treewidth-finiteness, coincides with a version of the notion of MSO-tree-interpretability, and is preserved under Courcelle-style MSO-transductions. We next sketch how the latter can used to cover DLs without the finite-treewidth model property.

Given an interpretation $\mathcal{I}$, we define *two-way regular path expressions* (2RPEs) as regular expressions over the alphabet containing the symbols $\mathtt{P}^{\rightarrow}$ and $\mathtt{P}^{\leftarrow}$ for binary predicates $\mathtt{P}$ from $\mathcal{I}$. Any 2RPE E defines a binary relation $E^{\mathcal{I}}$ by letting:

$$\begin{aligned}
(\mathtt{R}^{\rightarrow})^{\mathcal{I}} &= \mathtt{R}^{\mathcal{I}}\\
(\mathtt{R}^{\leftarrow})^{\mathcal{I}} &= \{(t_2,t_1) \mid (t_1,t_2) \in \mathtt{R}^{\mathcal{I}}\}\\
(E_1E_2)^{\mathcal{I}} &= \{(t_1,t_3) \mid (t_1,t_2) \in (E_1)^{\mathcal{I}}, (t_2,t_3) \in (E_2)^{\mathcal{I}}\}\\
(E_1 \cup E_2)^{\mathcal{I}} &= (E_1)^{\mathcal{I}} \cup (E_2)^{\mathcal{I}}\\
(E^*)^{\mathcal{I}} &= \{(t,t) \mid t \in \mathsf{dom}(\mathcal{I})\} \cup (E)^{\mathcal{I}} \cup (EE)^{\mathcal{I}} \cup \ldots
\end{aligned}$$

Now, for a binary predicate $\mathtt{P}$, we define the *2RPE-extension of* $\mathcal{I}$ w.r.t. $\mathtt{P}$ and E as the instance $\mathcal{I}^{\mathtt{P}:=\mathtt{P}\cup E}$ which maps $\mathtt{P}$ to $\mathtt{P}^{\mathcal{I}} \cup E^{\mathcal{I}}$ and behaves exctly like $\mathcal{I}$ otherwise.

Theorem 5. *If* $\mathcal{I}$ *has finite partitionwidth, then so does any 2RPE-extension of* $\mathcal{I}$*.*

Example 6. The description logics $\mathcal{SRIQ}b_s$, $\mathcal{SROI}b_s$, and $\mathcal{SROQ}b_s$ are extensions of the mentioned $\mathcal{ALCHIQ}b$, $\mathcal{ALCHOI}b$, and $\mathcal{ALCHOQ}b$, respectively. One additional feature of $\mathcal{SR}*$ knowledge bases is an *RBox*, which contains *role-chain axioms* like transitivity statements for binary relations such as $\mathtt{Friend} \circ \mathtt{Friend} \sqsubseteq \mathtt{Friend}$ but also more complex relationships like $\mathtt{Friend} \circ \mathtt{Enemy} \sqsubseteq \mathtt{Enemy}$. Such axioms prevent the existence of treewidth-finite (finitely) universal model sets, but not of partitionwidth-finite ones, as can be shown via known model-theoretic insights: Taking some model, it can be unravelled into a "pre-model" of finite treewidth (and thus partitionwidth), as one would do in the RBox-free case. This structure can then be extended to a model proper, coping with the still violated RBox axioms by applying RPE-extensions which – as we now know – preserve finiteness of partitionwidth. It is easy to see that the model thus obtained maps homomorphically into the original one. Hence, we obtain decidability of homomorphism-closed effectively MSO-expressible $\forall$SO queries for all these description logics. For query languages up to UC2RPQs, this had already been known (even with tight complexity bounds), but for more expressive query languages, decidability had hitherto been open.

[1] It should be noted that this corollary does not cover the case $\mathfrak{F} = \mathcal{ALCHOIQ}b$ from Example 2, since the employed construction to convert an arbitrary countermodel $\mathcal{I}$ into a finite-treewidth countermodel $\mathcal{J}$ does not ensure that $\mathcal{J}$ homomorphically maps into $\mathcal{I}$. In fact, the decidability of the entailment of queries beyond UCQs for $\mathcal{ALCHOIQ}b$ is wide open.

Contributions to Existential Rules. In our article, we also discuss how our decidability framework can be fruitfully applied to the very popular logical formalism of *existential rules* – also referred to as *tuple-generating dependencies* (TGDs), *conceptual graph rules*, Datalog$^{\pm}$, and $\forall\exists$*-rules* – which has become a de-facto standard of ontological querying over databases. As entailment of even the most simple queries is undecidable for unconstrained existential rules, there has been ample research into fragments that restore decidable querying. This includes investigations based on width notions, although existing work almost entirely focuses on (single) universal models of finite treewidth. We advance this line of research by introducing the notion of *finite partitionwidth sets* of existential rules, for which the existence of a universal model of finite partitionwidth is guaranteed even when combined with arbitrary finite sets of ground facts. We propose advanced stratification-based ways to analyse complex rule sets to find out if they fall into this category.

Conclusion. We are optimistic that the generic methods and results presented in this article can advance the understanding of general model-theoretic foundations of decidability of entailment problems. We expect that many more useful width notions can be defined (or discovered in the existing literature), leading to novel avenues for establishing decidability of expressive entailment problems. Thus, it seems worthwhile to investigate other (existing or newly designed) width notions regarding their friendliness properties. As less expressive $\mathfrak{L}$ may admit more general $\mathfrak{L}$-friendly w, there is a Pareto-optimality boundary to explore. For w-controllability of specification languages $\mathfrak{F}$ wrt. query languages $\mathfrak{Q}$, there is a three-way tradeoff between the generality of w, the expressivity of $\mathfrak{F}$, and the expressivity of $\mathfrak{Q}$.

Finally, as we have already demonstrated in this paper, we are confident that the width-based framework can guide the search for novel, very expressive, syntactically defined combinations of specification and query languages with a decidable entailment problem. Of course, for cases thus identified, subsequent investigations would have to determine the precise corresponding complexities.