

Extending NoHR for OWL 2 QL*

Nuno Costa, Matthias Knorr and João Leite

NOVA LINCS

Departamento de Informática

Universidade NOVA de Lisboa

2829-516 Caparica, Portugal

Abstract

The Protégé plug-in NoHR allows the user to combine an OWL 2 EL ontology with a set of non-monotonic (logic programming) rules – suitable, e.g., to express defaults and exceptions – and query the combined knowledge base (KB). The formal approach realized in NoHR is polynomial (w.r.t. data complexity) and it has been shown that even very large health care ontologies, such as SNOMED CT, can be handled. As each of the tractable OWL profiles is motivated by different application cases, extending the tool to the other profiles is of particular interest, also because these preserve the polynomial data complexity of the combined formalism. Yet, a straightforward adaptation of the existing approach to OWL 2 QL turns out to not be viable. In this paper, we provide the non-trivial solution for the extension of NoHR to OWL 2 QL by directly translating the ontology into rules without any prior pre-processing or classification. We have implemented our approach and our evaluation shows encouraging results.

1 Introduction

NoHR¹ is a plug-in for the ontology editor Protégé² that allows its users to query combinations of \mathcal{EL}_{\perp}^+ ontologies and non-monotonic rules in a top-down manner.

Its motivation stems from the fact that many current ontologies, such as the very large health care ontologies widely used in the area of medicine, e.g., SNOMED CT,³ are expressed in OWL 2 EL, one of the OWL 2 profiles [Motik *et al.*, 2013], and its underlying description logic (DL) \mathcal{EL}^{++} [Baader *et al.*, 2005]. Yet, due to their monotonic semantics, i.e., previously drawn conclusions persist when new additional information is adopted, DL-based ontology languages [Baader

et al., 2010] are not suitable to model defaults and exceptions with a closed-world view, a frequently requested feature, e.g., when matching patient records to clinical trial criteria [Patel *et al.*, 2007].

Among the plethora of approaches for extending DLs with non-monotonic features and deal with this problem (c.f. related work in [Eiter *et al.*, 2008; Motik and Rosati, 2010]), NoHR builds on Hybrid MKNF [Motik and Rosati, 2010], which is based on the logic of minimal knowledge and negation as failure (MKNF) [Lifschitz, 1991], under the well-founded semantics [Knorr *et al.*, 2011], a formalism that combines DLs and non-monotonic rules as known from Logic Programming (see also [Alberti *et al.*, 2012] for further motivation in its favor).

This choice is motivated, on the one hand, by the fact that non-monotonic logic programming rules are one of the most well-studied formalisms that admit expressing defaults, exceptions, and also integrity constraints in a declarative way, and are part of RIF [Boley and Kifer, 2013], the other expressive language for the Semantic Web whose standardization is driven by the W3C.⁴ On the other hand, Hybrid MKNF provides a very general and flexible framework for combining DL ontologies and non-monotonic rules (see [Motik and Rosati, 2010]). In addition, [Knorr *et al.*, 2011], which is a variant of [Motik and Rosati, 2010] based on the well-founded semantics [Gelder *et al.*, 1991] for logic programs, has a (lower) polynomial data complexity and is amenable for applying top-down query procedures, such as $\text{SLG}(\mathcal{O})$ [Alferes *et al.*, 2013], to answer queries based only on the information relevant for the query, and without computing the entire model.

NoHR is thus applicable to combinations of non-monotonic rules and OWL 2 EL ontologies. However, other applications (see, e.g., [Calvanese *et al.*, 2011; Savo *et al.*, 2010]) require ontologies using DL constructors which are not covered by OWL 2 EL, such as concept and role negation or role inverses, as admitting these would raise its polynomial complexity [Baader *et al.*, 2005].

OWL 2 QL and the *DL-Lite* family [Calvanese *et al.*, 2007; Artale *et al.*, 2009] to which the DL underneath OWL 2 QL belongs, *DL-Lite_R*, is suitable in these cases and has recently drawn a lot of attention in research and in applications.

*Partially supported by Fundação para a Ciência e a Tecnologia under project PTDC/EIA-CCO/121823/2010 and strategic project PEst/UID/CEC/04516/2013. M. Knorr was also supported by grant SFRH/BPD/86970/2012.

¹<http://centria.di.fct.unl.pt/nohr/>

²<http://protege.stanford.edu>

³<http://www.ihtsdo.org/snomed-ct/>

⁴<http://www.w3.org>

Even though a simple language at first glance, it is expressive enough to capture basic ontology languages, conceptual data models, e.g., Entity-Relationship, and object-oriented formalisms, e.g., basic UML class diagrams. Reasoning focuses on answering queries by rewriting the initial query, with the help of the ontology, into a set of queries that can be answered using an industry-strength SQL engine over the data. This provides the very low data complexity of LOGSPACE for query answering, but also links directly to applications in ontology-based data access (OBDA) [Calvanese *et al.*, 2011; Kontchakov *et al.*, 2011]. Altogether, OWL 2 QL is naturally tailored towards huge datasets.

In order to provide also such applications based on OWL 2 QL with the additional expressive power obtained from combining DL ontologies with non-monotonic rules, in this paper, we extend NoHR to OWL 2 QL. Whereas, at first sight, this could seem like a routine exercise, the fact that, to the best of our knowledge, no dedicated open-source OWL 2 QL classifier with OWL API is available, and applying the \mathcal{EL} reasoner ELK [Kazakov *et al.*, 2013], currently used in NoHR, to classify a $DL-Lite_R$ ontology is obviously not possible, we have to follow a different path here, namely translate the ontology directly into rules. This introduces some non-trivial problems, in particular, the need to capture unsatisfiable concepts and roles and irreflexive roles, for which in [Calvanese *et al.*, 2007] a closure of so-called negative axioms is computed, potentially introducing a huge number of additional axioms. We solve this problem by introducing an extension of the graph, used, e.g., in [Lembo *et al.*, 2013] for classification in OWL QL, to negative axioms. The resulting translation is implemented as a module of the NoHR translator, and its performance evaluated. Our main contributions are:

- A procedure for translating $DL-Lite_R$ ontologies into rules which allows answering queries over hybrid KBs combining such ontologies and non-monotonic rules;
- An substantial extension of the Protégé plug-in NoHR to include OWL 2 QL ontologies, beyond $DL-Lite_R$ via normalizations, including optimizations on the number of created rules and the use of tabling in the top-down query engine XSB Prolog;⁵
- An evaluation of our extension that shows that NoHR for OWL 2 QL maintains all positive evaluation results of the OWL 2 EL version [Ivanov *et al.*, 2013], and is even faster during pre-processing, as no classification is necessary, in exchange for an on average slightly longer response time during querying.

The remainder of the paper is structured as follows. In Sect. 2, we briefly recall $DL-Lite_R$ and MKNF knowledge bases as a tight combination of the former DL and non-monotonic rules. Then, we present the translation of $DL-Lite_R$ ontologies which allows us to query such MKNF knowledge bases in Sect. 3. In Sect. 4, we discuss the changes made in the implementation for OWL 2 QL including optimizations, and evaluate it in Sect. 5, before we conclude in Sect. 6.

⁵<http://xsb.sourceforge.net>

2 Preliminaries

2.1 $DL-Lite_{\mathcal{R}}$

The description logic underlying OWL QL is $DL-Lite_{\mathcal{R}}$, one language of the $DL-Lite$ family [Calvanese *et al.*, 2007; Artale *et al.*, 2009], which we recall following the presentation in [Knorr and Alferes, 2011].

The syntax of $DL-Lite_{\mathcal{R}}$ is based on three disjoint sets of *individual names* N_I , *concept names* N_C , and *role names* N_R . *Complex concepts* and *roles* can be formed according to the following grammar

$$B \rightarrow A \mid \exists Q \quad C \rightarrow B \mid \neg B \quad Q \rightarrow P \mid P^- \quad R \rightarrow Q \mid \neg Q$$

where $A \in N_C$ is a concept name, $P \in N_R$ a role name, and P^- its inverse. We also call B a *basic concept*, Q a *basic relation*, C a *general concept* and R a *general role*.

A $DL-Lite_{\mathcal{R}}$ knowledge base $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ consists of a TBox \mathcal{T} and an ABox \mathcal{A} . The TBox contains *general inclusion axioms (GCI)* of the form $B \sqsubseteq C$ and *role inclusion axioms (RI)* of the form $Q \sqsubseteq R$, with B, C, Q , and R defined as above. We term *positive inclusion axioms* all GCIs and RIs in \mathcal{O} such that C is a basic concept and R is a basic relation, respectively, and all other GCIs and RIs *negative inclusion axioms*. We also assume that Q^- denotes the role P if $Q = P^-$, and P^- if $Q = P$. The ABox contains assertions of the form $A(a)$ and $P(a, b)$ where $A \in N_C$, $P \in N_R$, and $a, b \in N_I$. Assertions $C(a)$ for general concepts C can be included by $A \sqsubseteq C$ and $A(a)$ for a new concept name A .

The semantics of $DL-Lite_{\mathcal{R}}$ is based on *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \mathcal{I}^{\mathcal{I}})$ consisting of a nonempty interpretation domain $\Delta^{\mathcal{I}}$ and an interpretation function $\mathcal{I}^{\mathcal{I}}$ that assigns to each individual a a distinct⁶ element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, to each concept name A a subset $A^{\mathcal{I}}$, and to each role name P a binary relation $P^{\mathcal{I}}$ over \mathcal{I} . This can be extended as usual:

$$\begin{aligned} (P^-)^{\mathcal{I}} &= \{(i_2, i_1) \mid (i_1, i_2) \in P^{\mathcal{I}}\} & (\neg B)^{\mathcal{I}} &= \Delta^{\mathcal{I}} \setminus B^{\mathcal{I}} \\ (\exists Q)^{\mathcal{I}} &= \{i \mid (i, i') \in Q^{\mathcal{I}}\} & (\neg Q)^{\mathcal{I}} &= \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \setminus Q^{\mathcal{I}} \end{aligned}$$

An interpretation \mathcal{I} is a *model of GCI* $B \sqsubseteq C$ and of *RI* $Q \sqsubseteq R$ if $B^{\mathcal{I}} \subseteq C^{\mathcal{I}}$ and $Q^{\mathcal{I}} \subseteq R^{\mathcal{I}}$ respectively. \mathcal{I} is also a *model of an assertion* $A(a)$ ($P(a, b)$) if $a^{\mathcal{I}} \in A^{\mathcal{I}}$ ($(a^{\mathcal{I}}, b^{\mathcal{I}}) \in P^{\mathcal{I}}$). Given an axiom/assertion α we denote by $\mathcal{I} \models \alpha$ that \mathcal{I} is a model of α . A *model of a $DL-Lite_{\mathcal{R}}$ KB* $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ is an interpretation \mathcal{I} such that $\mathcal{I} \models \alpha$ holds for all $\alpha \in \mathcal{T} \cup \mathcal{A}$, and \mathcal{O} is *satisfiable* if it has at least one model, and *unsatisfiable* otherwise. Also, \mathcal{O} *entails* axiom α , written $\mathcal{O} \models \alpha$, if every model of \mathcal{O} satisfies α .

2.2 MKNF Knowledge Bases

MKNF knowledge bases (KBs) build on the logic of minimal knowledge and negation as failure (MKNF) [Lifschitz, 1991]. Two main different semantics have been defined [Motik and Rosati, 2010; Knorr *et al.*, 2011], and we focus on the well-founded version [Knorr *et al.*, 2011], due to its lower computational complexity and amenability to top-down querying without computing the entire model. Here, we only point out

⁶Hence, the unique name assumption is applied and, as shown in [Artale *et al.*, 2009], dropping it would increase significantly the computational complexity of $DL-Lite_{\mathcal{R}}$.

important notions following [Ivanov *et al.*, 2013], and refer to [Knorr *et al.*, 2011] and [Alferes *et al.*, 2013] for the details.

We start by recalling MKNF knowledge bases as presented in [Alferes *et al.*, 2013] to combine an ontology and a set of non-monotonic rules (similar to a normal logic program).

Definition 1. Let \mathcal{O} be an ontology. A function-free first-order atom $P(t_1, \dots, t_n)$ s.t. P occurs in \mathcal{O} is called *DL-atom*; otherwise *non-DL-atom*. A rule r is of the form

$$H \leftarrow A_1, \dots, A_n, \text{not}B_1, \dots, \text{not}B_m \quad (1)$$

where the *head* of r , H , and all A_i with $1 \leq i \leq n$ and B_j with $1 \leq j \leq m$ in the *body* of r are atoms. A *program* \mathcal{P} is a finite set of rules, and an *MKNF knowledge base* \mathcal{K} is a pair $(\mathcal{O}, \mathcal{P})$. A rule r is *DL-safe* if all its variables occur in at least one non-DL-atom A_i with $1 \leq i \leq n$, and \mathcal{K} is *DL-safe* if all its rules are DL-safe.

DL-safety ensures decidability of reasoning with MKNF knowledge bases and can be achieved by introducing a new predicate o , adding $o(i)$ to \mathcal{P} for all constants i appearing in \mathcal{K} and, for each rule $r \in \mathcal{P}$, adding $o(X)$ for each variable X appearing in r to the body of r . Therefore, we only consider DL-safe MKNF knowledge bases.

Example 1. Consider an MKNF knowledge base \mathcal{K} as given below for recommending CDs adapted from [Knorr *et al.*, 2011] (with some modifications). We denote DL-atoms and constants with upper-case names and non-DL-atoms and variables with lower-case names.⁷

$$\begin{aligned} \exists \text{HasArtist}^- &\sqsubseteq \text{Artist} & \text{Piece} &\sqsubseteq \exists \text{HasArtist} \\ \exists \text{HasComposed}^- &\sqsubseteq \text{Piece} & \text{Artist} &\sqsubseteq \neg \text{Piece} \\ \text{HasComposed}^- &\sqsubseteq \text{HasArtist} \\ \text{recommend}(x) &\leftarrow \text{Piece}(x), \text{not} \text{owns}(x), \text{not} \text{lowEval}(x), \\ &\text{interesting}(x) \\ \text{interesting}(x) &\leftarrow \text{Piece}(x), \text{not} \text{owns}(x), \text{Piece}(y), \text{owns}(y), \\ &\text{Artist}(z), \text{HasArtist}(y, z), \text{HasArtist}(x, z) \\ \text{owns}(\text{Summertime}) &\leftarrow \text{Piece}(\text{Summertime}) \leftarrow \\ &\text{HasArtist}(\text{Summertime}, \text{Gershwin}) \leftarrow \\ &\text{HasComposed}(\text{Gershwin}, \text{RhapsodyInBlue}) \leftarrow \end{aligned}$$

This example shows that we can seamlessly express defaults and exceptions, such as recommending pieces as long as they are not owned or having a low evaluation, and at the same time taxonomic/ontological knowledge including information over unknown individuals, such as every piece having at least one artist without having to specify whom, but also features of *DL-Lite_R*, such as domain and range restrictions (of roles).

The semantics of MKNF knowledge bases \mathcal{K} is usually given by a translation π into an MKNF formula $\pi(\mathcal{K})$, i.e., a formula over first-order logic extended with two modal operators \mathbf{K} and not . Namely, every rule of the form (1) is translated into $\mathbf{K}H \leftarrow \mathbf{K}A_1, \dots, \mathbf{K}A_n, \text{not}B_1, \dots, \text{not}B_m$, $\pi(\mathcal{P})$ is the conjunction of the translations of its rules, and $\pi(\mathcal{K}) = \mathbf{K}\pi(\mathcal{O}) \wedge \pi(\mathcal{P})$ where $\pi(\mathcal{O})$ is the first-order translation of \mathcal{O} . Reasoning with such MKNF formulas is then

⁷To ease readability, we omit the auxiliary atoms that ensure DL-safety and leave them implicit.

commonly achieved using a partition of *modal atoms*, i.e., all expressions of the form $\mathbf{K}\varphi$ for each $\mathbf{K}\varphi$ or $\text{not}\varphi$ occurring in $\pi(\mathcal{K})$. For [Knorr *et al.*, 2011], such a partition assigns *true*, *false*, or *undefined* to (modal) atoms, and can be effectively computed in polynomial time. If \mathcal{K} is *MKNF-consistent*, then this partition does correspond to the unique model of \mathcal{K} [Knorr *et al.*, 2011], and, like in [Alferes *et al.*, 2013], we call the partition the *well-founded MKNF model* $M_{\text{wf}}(\mathcal{K})$. Here, \mathcal{K} may indeed not be MKNF-consistent if the ontology alone is unsatisfiable, or by the combination of appropriate axioms in \mathcal{O} and rules in \mathcal{P} , e.g., $A \sqsubseteq \neg B$, and $A(a) \leftarrow$ and $B(a) \leftarrow$. Strictly speaking, unlike [Ivanov *et al.*, 2013], we do not have to make assumptions on the satisfiability of \mathcal{O} as we are not going to use a classifier when processing *DL-Lite_R* ontologies. Still, for the technical results established in Sec. 3, we will rely on satisfiability, since we are able to entail everything from an unsatisfiable \mathcal{O} , whereas the translation into rules defined in Sec. 3 would not permit that. This is why in the following, we assume that \mathcal{O} occurring in \mathcal{K} is satisfiable, which does not truly constitute a restriction as we can always turn the ABox into rules without any effect on $M_{\text{wf}}(\mathcal{K})$. An alternative approach would be to use one of the paraconsistent semantics for MKNF knowledge bases [Kaminski *et al.*, 2015], but this is outside the scope of this paper, and an issue for future work as currently no paraconsistent correspondence to the querying procedure $\text{SLG}(\mathcal{O})$ used here exists.

2.3 Querying in MKNF Knowledge Bases

In [Alferes *et al.*, 2013], a procedure, called $\text{SLG}(\mathcal{O})$, is defined for querying MKNF knowledge bases under the well-founded MKNF semantics. This procedure extends SLG resolution with tabling [Chen and Warren, 1996] with an *oracle* to \mathcal{O} that handles ground queries to the DL-part of \mathcal{K} by returning (possibly empty) sets of atoms that, together with \mathcal{O} and information already proven true, allows us to derive the queried atom. We refer to [Alferes *et al.*, 2013] for the full account of $\text{SLG}(\mathcal{O})$, and only recall a few crucial notions necessary in the following.

$\text{SLG}(\mathcal{O})$ is based on creating top-down derivation trees with the aim of answering (*DL-safe*) conjunctive queries $Q = q(\vec{X}) \leftarrow A_1, \dots, A_n, \text{not}B_1, \dots, \text{not}B_m$ where each variable in Q occurs in at least one non-DL atom in Q , and where \vec{X} is the (possibly empty) set of requested variables appearing in the body.

In general, the computation of $M_{\text{wf}}(\mathcal{K})$ uses two different versions of \mathcal{K} in parallel to guarantee that a) coherence is ensured, i.e., if $\neg P(a)$ is derivable, then $\text{not}P(a)$ has to be true as well (cf. also [Knorr *et al.*, 2011]), and b) MKNF-consistency of \mathcal{K} can be verified. For a top-down approach this is impractical, so, instead, a doubled MKNF knowledge base $\mathcal{K}^d = (\mathcal{O}, \mathcal{O}^d, \mathcal{P}^d)$ is defined in which a copy of \mathcal{O} with new doubled predicates is added, and two rules occur in \mathcal{P}^d for each rule in \mathcal{P} , intertwining original and doubled predicates (see Def. 3.1 in [Alferes *et al.*, 2013]). It is shown that an atom A is true in $M_{\text{wf}}(\mathcal{K})$ iff A is true in $M_{\text{wf}}(\mathcal{K}^d)$ and A is false in $M_{\text{wf}}(\mathcal{K})$ iff A^d is false in $M_{\text{wf}}(\mathcal{K}^d)$. Note that \mathcal{K}^d is necessary in general, but we can use \mathcal{K} here if it contains no negative inclusion axioms.

In [Alferes *et al.*, 2013], the notion of oracle is defined to handle ground queries to the ontology, but before we recall that notion, we use an example to illustrate the idea.

Example 2. Recall \mathcal{K} in Ex. 1. Here, we omit \mathcal{K}^d and restrict ourselves to \mathcal{K} , which suffices our purposes. Consider query $q = \text{recommend}(\text{Summertime})$. By instantiating the body of the matching rule head in \mathcal{K} with $x = \text{Summertime}$, we obtain two new queries. The first one, $\text{Piece}(\text{Summertime})$, can be answered by means of the rule with matching head. The second, $\text{notowns}(\text{Summertime})$, is handled by querying for $\text{owns}(\text{Summertime})$, for which a corresponding rule exists, so $\text{notowns}(\text{Summertime})$ fails, hence q is false.

Consider $q_1 = \text{recommend}(\text{RhapsodyInBlue})$. Using the same rule with matching rule head we obtain four new instantiated queries from the rule body. Now, $\text{Piece}(\text{RhapsodyInBlue})$ cannot be derived from the rules, but we can query the ontology and the oracle will return, e.g., a query $\text{HasComposed}(x_1, \text{RhapsodyInBlue})$ that if proven true can be added to \mathcal{O} , which would allow us to derive the queried goal. This query succeeds because of $\text{HasComposed}(\text{Gershwin}, \text{RhapsodyInBlue}) \leftarrow$, and so does $\text{Piece}(\text{RhapsodyInBlue})$. Then, we cannot prove $\text{owns}(\text{RhapsodyInBlue})$ nor $\text{lowEval}(\text{RhapsodyInBlue})$, so both fail, succeeding their (default) negated queries. For the remaining new query $\text{interesting}(\text{RhapsodyInBlue})$, the second rule head matches, creating further subgoals. The first two were just answered, as the next two with $y = \text{Summertime}$ for q . The remaining also follow from the interplay of \mathcal{O} and \mathcal{P} in \mathcal{K} , so q_1 succeeds.

We recall the notions of a complete and a (correct) partial oracle from [Alferes *et al.*, 2013].

Definition 2. Let $\mathcal{K}^d = (\mathcal{O}, \mathcal{O}^d, \mathcal{P}^d)$ be a doubled MKNF KB, \mathcal{I} a set of ground atoms (already proven to be true), S a ground query, and \mathcal{L} a set of ground atoms such that each $L \in \mathcal{L}$ is unifiable with at least one rule head in \mathcal{P}^d . The *complete oracle* for \mathcal{O} , denoted $\text{comp}T_{\mathcal{O}}$, is defined by $\text{comp}T_{\mathcal{O}}(\mathcal{I}, S, \mathcal{L})$ iff $\mathcal{O} \cup \mathcal{I} \cup \mathcal{L} \models S$ or $\mathcal{O}^d \cup \mathcal{I} \cup \mathcal{L} \models S$. A *partial oracle* for \mathcal{O} , denoted $pT_{\mathcal{O}}$, is a relation $pT_{\mathcal{O}}(\mathcal{I}, S, \mathcal{L})$ such that if $pT_{\mathcal{O}}(\mathcal{I}, S, \mathcal{L})$, then $\mathcal{O} \cup \mathcal{I} \cup \mathcal{L} \models S$ or $\mathcal{O}^d \cup \mathcal{I} \cup \mathcal{L} \models S$ for consistent $\mathcal{O} \cup \mathcal{I} \cup \mathcal{L}$ and $\mathcal{O}^d \cup \mathcal{I} \cup \mathcal{L}$, respectively.

A partial oracle $pT_{\mathcal{O}}$ is *correct* w.r.t. $\text{comp}T_{\mathcal{O}}$ iff, for all MKNF-consistent \mathcal{K}^d , replacing $\text{comp}T_{\mathcal{O}}$ in $\text{SLG}(\mathcal{O})$ with $pT_{\mathcal{O}}$ succeeds for exactly the same set of queries.

Partial oracles may avoid returning unnecessary answers \mathcal{L} , such as non-minimal answers or those that try to derive an MKNF-inconsistency even though \mathcal{K}^d is MKNF-consistent. Also, correctness of partial oracles is only defined w.r.t MKNF-consistent \mathcal{K} . The rationale is that, when querying top-down, we want to avoid checking whether the entire KB \mathcal{K}^d is MKNF-consistent. This leads to para-consistent derivations if \mathcal{K}^d is not MKNF-consistent, e.g., some atom P is true, yet P^d is false, while other independent atoms are evaluated as if \mathcal{K}^d was MKNF-consistent (see [Alferes *et al.*, 2013]).

3 Translating the Ontology into Rules

As argued for the case of \mathcal{EL}_{\perp}^+ [Ivanov *et al.*, 2013], axioms with \exists on the right-hand side, e.g., $\text{Piece} \sqsubseteq \exists \text{HasArtist}$,

cannot be translated straightforwardly into rules, nor do they directly contribute to the result when querying for ground instances, e.g., of $\text{HasArtist}(x, y)$. Still, such axioms may contribute to derivations within \mathcal{O} , which is why, in [Ivanov *et al.*, 2013], a classification using the dedicated and highly efficient \mathcal{EL} reasoner ELK [Kazakov *et al.*, 2013] is first applied to derive implicit consequences. These, together with all axioms in \mathcal{O} , are then translated into rules, now discarding certain axioms with \exists on the right-hand side.

Here, since to the best of our knowledge no dedicated, open-source OWL 2 QL classifier with OWL API is available, we opt to follow a different path, namely translate the ontology directly into rules. This also simplifies and shortens the preprocessing phase and avoids a priori-classification, but requires some non-trivial considerations to ensure that no derivations are lost in the process, which we will explain next.

Essentially, axioms, such as $\text{Piece} \sqsubseteq \exists \text{HasArtist}$, cannot be translated into a rule $\text{HasArtist}(x, y) \leftarrow \text{Piece}(x)$ using a universal variable y , as this would allow us to derive $\text{HasArtist}(x, y)$ for any $\text{Piece}(x)$ and y , which is clearly not what the axiom expresses. Using a new constant c instead of y would not be correct either, as querying for $\text{HasArtist}(x, y)$ would return $\text{HasArtist}(x, c)$ for any $\text{Piece}(x)$ for the same c . Therefore, we proceed differently by introducing new auxiliary predicates that intuitively represent the domain and range of roles. For our example, this will yield the rule $D\text{HasArtist}(x) \leftarrow \text{Piece}(x)$ where $D\text{HasArtist}$ stands for the domain of HasArtist (and $R\text{HasArtist}$ its range). Using such auxiliary predicates also means that we have to make sure that, e.g., $\text{HasArtist}(\text{Summertime}, \text{Gershwin})$ allows us to derive $D\text{HasArtist}(\text{Summertime})$, which can be achieved via an additional rule $D\text{HasArtist}(x) \leftarrow \text{HasArtist}(x, y)$. Moreover, for $\text{HasComposed}^- \sqsubseteq \text{HasArtist}$, it does not suffice to translate the axiom to $\text{HasArtist}(x, y) \leftarrow \text{HasComposed}(y, x)$, but also link the new auxiliary predicates for both roles, by adding, $D\text{HasArtist}(x) \leftarrow R\text{HasComposed}(x)$ and $R\text{HasArtist}(x) \leftarrow D\text{HasComposed}(x)$.

We now formalize this translation, and we start by introducing notation on how to translate general concepts and roles. For that purpose, we formally introduce for each role $P \in \mathbb{N}_{\mathbb{R}}$ auxiliary predicates DP and RP with the intuition of representing the domain and range of P . Also, similar to previous work in [Alferes *et al.*, 2013; Ivanov *et al.*, 2013], we use special atoms $NH(\vec{t}_i)$ in $\text{SLG}(\mathcal{O})$ that represent a query $\neg H(\vec{t}_i)$ to the oracle. These are, of course, only relevant if \mathcal{O} contains negative inclusion axioms.

Definition 3. Let C be a concept, R a role, x and y variables, and v a new (anonymous) variable (disjoint from x and y).

We define $tr(C, x)$ and $tr(R, x, y)$ as follows:

$$tr(C, x) = \begin{cases} A(x) & \text{if } C = A \\ DP(x) & \text{if } C = \exists P \\ RP(x) & \text{if } C = \exists P^- \\ NA(x) & \text{if } C = \neg A \\ tr(\neg Q, x, v) & \text{if } C = \neg \exists Q \end{cases}$$

$$tr(R, x, y) = \begin{cases} P(x, y) & \text{if } R = P \\ P(y, x) & \text{if } R = P^- \\ NP(x, y) & \text{if } R = \neg P \\ NP(y, x) & \text{if } R = \neg P^- \end{cases}$$

We obtain $tr^d(C, x)$ and $tr^d(Q, x, y)$ from $tr(C, x)$ and $tr(Q, x, y)$ by substituting all predicates P in $tr(C, x)$ and $tr(Q, x, y)$ with P^d , respectively.

$tr(C, x)$ and $tr(R, x, y)$ handle both positive and negative inclusions and no additional case distinction is necessary.

Before we present the actual translation, we need to introduce one central notion, namely a graph to represent the axioms in a given TBox \mathcal{T} as well as the implicitly derivable axioms, which will be necessary for defining the translation itself, but also turn out useful when establishing the correctness of the translation. Graphs have been used for classification in OWL QL (of positive inclusion axioms) [Lembo *et al.*, 2013], and we extend the notion here to also take negative inclusion axioms into account. We thus introduce the digraph (directed graph) of \mathcal{T} as follows.

Definition 4. Let \mathcal{T} be a $DL\text{-}Lite_R$ TBox. The *digraph* of \mathcal{T} , $\mathcal{G}_{\mathcal{T}} = \langle \mathcal{V}, \mathcal{E} \rangle$, is constructively defined as follows.

1. If $A \in \mathbb{N}_C$, then A and $\neg A$ are in \mathcal{V} ;
2. If $R \in \mathbb{N}_R$, then $P, \exists P, \exists P^-, \neg P, \neg \exists P$, and $\neg P^-$ are in \mathcal{V} ;
3. If $B_1 \sqsubseteq B_2 \in \mathcal{T}$, then the edges (B_1, B_2) and $(\neg B_2, \neg B_1)$ are in \mathcal{E} ;
4. If $Q_1 \sqsubseteq Q_2 \in \mathcal{T}$, then the edges $(Q_1, Q_2), (Q_1^-, Q_2^-), (\exists Q_1, \exists Q_2), (\exists Q_1^-, \exists Q_2^-), (\neg Q_2, \neg Q_1), (\neg Q_2^-, \neg Q_1^-), (\neg \exists Q_2, \neg \exists Q_1)$ and $(\neg \exists Q_2^-, \neg \exists Q_1^-)$ are in \mathcal{E} ;
5. If $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, then the edges $(B_1, \neg B_2)$ and $(B_2, \neg B_1)$ are in \mathcal{E} ;
6. If $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{T}$, then the edges $(Q_1, \neg Q_2), (Q_2^-, \neg Q_1^-), (\exists Q_1, \neg \exists Q_2), (\exists Q_2, \neg \exists Q_1), (\exists Q_1^-, \neg \exists Q_2^-)$ and $(\exists Q_2^-, \neg \exists Q_1^-)$ are in \mathcal{E} .

Basically, each possible general concept and general role over \mathbb{N}_C and \mathbb{N}_R is a node in $\mathcal{G}_{\mathcal{T}}$, and the directed edges represent logical implications that follow from the axioms. Namely, for items 3. and 5., the subset inclusion itself and its contrapositive are in \mathcal{E} , and this is similar for items 4. and 6., only that the additional combinations due to inverses, \exists , and \neg have to be taken into account. In this sense, the graph can be understood as capturing all subset inclusions (explicit and implicit) in \mathcal{O} , i.e., whenever there is a path from concept C_1 to concept C_2 and from role R_1 to role R_2 , then $C_1 \sqsubseteq C_2$ and $R_1 \sqsubseteq R_2$ hold respectively. An Example of such a digraph is given in Fig. 1 for the TBox \mathcal{T} from Example 1.

One observation to be made in Fig. 1, is that $\exists HasComposed \sqsubseteq \neg \exists HasComposed$, i.e., $HasComposed$

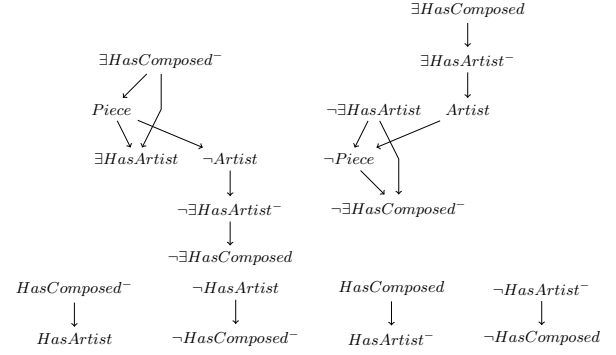


Figure 1: The digraph $\mathcal{G}_{\mathcal{T}}$ for Example 1

is irreflexive. Even though this does not entail any assertion, knowing that $\forall x. \neg HasComposed(x, x)$ does hold should be captured in the translation. We introduce $\Psi(\mathcal{T})$, the set of irreflexive roles in \mathcal{T} , to be able to ensure exactly that.

Definition 5. Let \mathcal{T} be a $DL\text{-}Lite_R$ TBox and $\mathcal{G}_{\mathcal{T}}$ its digraph. We define $\Psi(\mathcal{T})$ as the smallest set of all $P \in \mathbb{N}_R$ that satisfy at least one of the following conditions:

1. For some $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, there exist paths from $\exists P$ to B_1 and from $\exists P^-$ to B_2 ;
2. For some $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, there exist paths from $\exists P^-$ to B_1 and from $\exists P$ to B_2 ;
3. For some $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{T}$, there exist paths from P to Q_1 and from P^- to Q_2 ;
4. For some $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{T}$, there exist paths from P^- to Q_1 and from P to Q_2 .

This notion builds on $\mathcal{G}_{\mathcal{T}}$, which is also required for detecting a further set of derivations. Imagine we would (wrongfully) add $Artist \sqsubseteq \exists HasComposed^-$ to \mathcal{O} in Example 1. Then there would be a path from $Artist$ to both $Piece$ and $\neg Piece$, i.e., the concept $Artist$ would be unsatisfiable. Note that independently of whether the hybrid KB is MKNF-inconsistent or not, we need to make sure that all unsatisfiable concepts and roles are determined, so we introduce $\Omega(\mathcal{T})$, quite similar in spirit to $\Psi(\mathcal{T})$.

Definition 6. Let \mathcal{T} be a $DL\text{-}Lite_R$ TBox and $\mathcal{G}_{\mathcal{T}}$ its digraph. We define $\Omega(\mathcal{T})$ as the smallest set of all $A \in \mathbb{N}_C$ such that, for some $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, there exist paths from A to both B_1 and B_2 , and all $P \in \mathbb{N}_R$ that satisfy at least one of the following conditions:

1. For some $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, there exist paths from $\exists P$ to both B_1 and B_2 ;
2. For some $B_1 \sqsubseteq \neg B_2 \in \mathcal{T}$, there exist paths from $\exists P^-$ to both B_1 and B_2 ;
3. For some $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{T}$, there exist paths from P to both Q_1 and Q_2 ;
4. For some $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{T}$, there exist paths from P^- to both Q_1 and Q_2 .

With all pieces in place, we can finally introduce the definition of the translation of a $DL\text{-}Lite_R$ ontology into rules.

Definition 7. Let \mathcal{O} be a $DL\text{-}Lite_R$ ontology. We define $\mathcal{P}_{\mathcal{O}}^d$ from \mathcal{O} , where B_1, B_2 are basic concepts, Q_1, Q_2 basic

roles, x, y variables, and a, b individuals, as the smallest set containing:

- (e) for every $P \in \mathbf{N}_R$:
 $DP(x) \leftarrow P(x, y) \quad DP^d(x) \leftarrow P^d(x, y)$
 $RP(y) \leftarrow P(x, y) \quad RP^d(y) \leftarrow P^d(x, y)$
- (a1) for every $A(a) \in \mathcal{O}$:
 $A(a) \leftarrow A^d(a) \leftarrow \text{not}NA(a)$
- (a2) for every $P(a, b) \in \mathcal{O}$:
 $P(a, b) \leftarrow P^d(a, b) \leftarrow \text{not}NP(a, b)$
- (s1) for every $B_1 \sqsubseteq B_2 \in \mathcal{O}$:
 $tr(B_2, x) \leftarrow tr(B_1, x)$
 $tr^d(B_2, x) \leftarrow tr^d(B_1, x), \text{not}tr(\neg B_2, x)$
 $tr(\neg B_1, x) \leftarrow tr(\neg B_2, x)$
- (s2) for every $Q_1 \sqsubseteq Q_2 \in \mathcal{O}$:
 $tr(Q_2, x, y) \leftarrow tr(Q_1, x, y)$
 $tr^d(Q_2, x, y) \leftarrow tr^d(Q_1, x, y), \text{not}tr(\neg Q_2, x, y)$
 $tr(\exists Q_2, x) \leftarrow tr(\exists Q_1, x)$
 $tr^d(\exists Q_2, x) \leftarrow tr^d(\exists Q_1, x), \text{not}tr(\neg \exists Q_2, x)$
 $tr(\exists Q_2^-, x) \leftarrow tr(\exists Q_1^-, x)$
 $tr^d(\exists Q_2^-, x) \leftarrow tr^d(\exists Q_1^-, x), \text{not}tr(\neg \exists Q_2^-, x)$
 $tr(\neg Q_1, x, y) \leftarrow tr(\neg Q_2, x, y)$
- (n1) for every $B_1 \sqsubseteq \neg B_2 \in \mathcal{O}$:
 $tr(\neg B_1, x) \leftarrow tr(B_2, x) \quad tr(\neg B_2, x) \leftarrow tr(B_1, x)$
- (n2) for every $Q_1 \sqsubseteq \neg Q_2 \in \mathcal{O}$:
 $tr(\neg Q_2, x, y) \leftarrow tr(Q_1, x, y)$
 $tr(\neg Q_1, x, y) \leftarrow tr(Q_2, x, y)$
- (i1) for every $A \in \Omega(\mathcal{T})$: $NA(x) \leftarrow$
- (i2) for every $P \in \Omega(\mathcal{T})$: $NP(x, y) \leftarrow$
- (ir) for every $P \in \Psi(\mathcal{T})$: $NP(x, x) \leftarrow$

Item (e) ensures that the domain and range of roles is correctly encoded, items (a1) and (a2) translate the ABox, items (s1) and (s2) the positive inclusions, items (n1) and (n2) the negative inclusions, and items (i1), (i2), and (ir) introduce the rules representing unsatisfiable concepts and unsatisfiable and irreflexive roles. Note, that $\mathcal{P}_{\mathcal{O}}^d$ contains the rule representation for both \mathcal{O} and \mathcal{O}^d , which is why items (e)–(s2) contain doubled rules. Of course, if \mathcal{O} does not contain negative inclusion axioms, then we can skip all these, as well as items (n1)–(ir) which will not contribute anything anyway in this case. The additional default atoms are added to the doubled rules to be in line with the idea of the doubling of rules in [Alferes *et al.*, 2013]: whenever, e.g., $A(x)$ is “classically false” for some x , i.e., $NA(x)$ holds, then we make sure that $A^d(x)$ is derivable as false for that same x from the rules, but not necessarily $A(x)$, thus allowing to detect potential MKNF-inconsistencies. That is also the reason why neither (n1)–(ir) nor the contrapositives in (s1) and (s2) do produce the doubled counterparts: atoms based on predicates of the forms NC^d or NR^d are not used anywhere. Finally, the doubled rules in (e) do not contain the default negated atom as this case only associates domain and range to a role assertion, either present in the ABox or derived elsewhere. Additionally, predicates NDP or NRP are not used anywhere, so such default negated atoms would be of no impact.

We can establish three correspondences between entailment from satisfiable \mathcal{O} and the program resulting from the translation $\mathcal{P}_{\mathcal{O}}^d$. First, we consider positive atoms.

Lemma 1. *Let \mathcal{O} be a $DL\text{-Lite}_R$ ontology, A a unary and R a binary predicate:*

- $\mathcal{O} \models A(a)$ iff $\mathcal{P}_{\mathcal{O}}^d \models A(a)$ and $\mathcal{O} \models R(a, b)$ iff $\mathcal{P}_{\mathcal{O}}^d \models R(a, b)$.

A similar property holds for (classically) atoms.

Lemma 2. *Let \mathcal{O} be a $DL\text{-Lite}_R$ ontology, A a unary and R a binary predicate:*

- $\mathcal{O} \models \neg A(a)$ iff $\mathcal{P}_{\mathcal{O}}^d \models NA(a)$ and $\mathcal{O} \models \neg R(a, b)$ iff $\mathcal{P}_{\mathcal{O}}^d \models NR(a, b)$.

We can also show the correspondent to Lemma 1 for the doubled predicates.

Lemma 3. *Let \mathcal{O} be a $DL\text{-Lite}_R$ ontology, A a unary and R a binary predicate:*

- $\mathcal{O}^d \models A^d(a)$ iff $\mathcal{P}_{\mathcal{O}}^d \models A^d(a)$ and $\mathcal{O}^d \models R^d(a, b)$ iff $\mathcal{P}_{\mathcal{O}}^d \models R^d(a, b)$.

Thus, we can define a correct partial oracle based on $\mathcal{P}_{\mathcal{O}}^d$.

Theorem 4. *Let $\mathcal{K}^d = (\mathcal{O}, \mathcal{O}^d, \mathcal{P}^d)$ be a doubled MKNF KB and $pT_{\mathcal{O}}^{QL}$ a partial QL oracle such that $pT_{\mathcal{O}}^{QL}(\mathcal{I}, S, \mathcal{L})$ iff $\mathcal{P}_{\mathcal{O}}^d \cup \mathcal{I} \cup \mathcal{L} \models S$. Then $pT_{\mathcal{O}}^{QL}$ is a correct partial oracle w.r.t. $\text{comp}T_{\mathcal{O}}$.*

Instead of coupling two rule reasoners that interact with each other using an oracle, we can simplify the process altogether and integrate both into one rule reasoner. The resulting approach is decidable with polynomial data complexity.

Theorem 5. *Let $\mathcal{K} = (\mathcal{O}, \mathcal{P})$ be an MKNF KB with \mathcal{O} in $DL\text{-Lite}_R$. An $\text{SLG}(\mathcal{O})$ evaluation of a query in $\mathcal{K}_{QL} = (\emptyset, (\mathcal{P}^d \cup \mathcal{P}_{\mathcal{O}}^d))$ is decidable with data complexity in $P\text{TIME}$.*

4 System Description

In this section, we briefly describe the changes to the architecture of our plug-in and discuss some optimizations implemented w.r.t. the translation described in Sec. 3.

To allow the usage of OWL QL ontologies, changes were essentially made in the translator. First, since now two OWL profiles are supported we have introduced a switch that checks the profile of the loaded/edited ontology. If it is in OWL EL, then NoHR behaves as described in [Ivanov *et al.*, 2013], i.e., the reasoner ELK is used to classify the ontology and return the inferred axioms to translator, which are then translated. Otherwise, we treat \mathcal{O} of the hybrid KB based on the translation described in Sec. 3 for OWL QL.

Notably, in Sec. 3, we only considered $DL\text{-Lite}_R$ while OWL QL includes a number of additional constructs which often can be expressed in $DL\text{-Lite}_R$. To account for that, we first normalize such expressions to axioms in $DL\text{-Lite}_R$. This includes ignoring certain expressions, most of which do not contribute anything to derivations, e.g., `SubClassOf(B owl:Thing)`, while others make the ontology unsatisfiable, such as `ClassAssertion(owl:Nothing a)`, although, as mentioned before, with no effect when querying the translated rules. The details on the normalization can be found in the appendix of the extended paper.

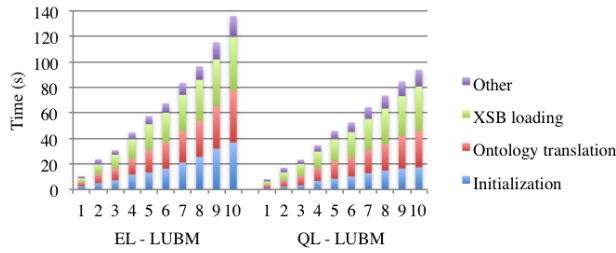


Figure 2: Preprocessing time for LUBM for the two translation modes

Subsequently, the graph is constructed, for determining unsatisfiable concepts and unsatisfiable and irreflexive roles, after which the translation is performed, which includes a number of optimizations. First, whenever there are no negative inclusions, the doubled rules are omitted in the cases (e)–(s2) of Def. 7. Additionally, case (e) is limited to those rules whose heads appear in the body of another rule. Both steps reduce the overall number of rules created during the translation.

The second group of optimizations is related to tabling in XSB, which contributes to help answering queries very efficiently in a top-down manner, and avoid infinite loops while querying. However, simply declaring all predicates to be tabled is very memory-consuming, so we reduced the number of tabled predicates without affecting loop detection. For example, only predicates that appear in any rule head and in any rule body need to be tabled. In addition, rules with an empty body (facts) can be ignored in the previous criterion, as these will not cause an infinite loop.

5 Evaluation

In this section, we evaluate our system and show that a) preprocessing is even faster when compared to NoHRs EL version, which was already capable of preprocessing large ontologies in a short period of time, b) querying scales well, even for over a million facts/assertions in the ABox, despite being slightly slower on average in comparison to EL, and c) adding rules scales linearly for pre-processing and querying, even for an ontology with many negative inclusions.

Tests were performed on a Notebook running Linux 3.17.6-1-ARCH (x86_64) with 1.8 GHz 4x Intel Core i3 processor and 4 GB of RAM. We used XSB 3.4.0 for querying, ran all tests in a terminal version and Java with “-XX:+AggressiveHeap” option, and report averages over 5 runs.

First, we considered LUBM⁸ [Guo *et al.*, 2005], a standard benchmark for evaluating queries over a large data set. The ontology itself is already rather simple and we reduced it even a bit further by removing all axioms that are not common to both OWL 2 QL and EL. The resulting ontology has only ninety logical axioms, but this way we can use it with both translators included in NoHR and compare their performance. We created instances of LUBM 1–10 with assertions ranging

⁸<http://swat.cse.lehigh.edu/projects/lubm/>




Figure 3: Query time for three LUBM queries

from roughly 100,000 to over 1,300,000 and performed preprocessing from loading the ontology to loading the translation result in XSB. The results for both translators EL and QL can be found in Fig. 2. Note that the segment “Initialization” is the time for preparing the translation, which for EL includes classifying the ontology, while the larger part of the segment “Other” corresponds to loading the ontology.

We can observe that QL is considerably faster, indeed up to 40s for LUBM10, to a considerable extent due to avoiding classification. Besides that, the preprocessing time increases linearly, and the overall time for preprocessing is acceptable in our opinion as this is only done once before querying.

Next, we also queried the resulting ten rule sets in XSB for both EL and QL using queries from the LUBM benchmark, that were manually transformed from SPARQL to queries usable in XSB. Among the fourteen provided queries, we chose seven, because the others were either no longer meaningful due to removal of certain axioms/DL constructors during the initial simplifications we applied to LUBM, or because initial tests revealed that XSB would run out of memory for a query, simply because, for our test system with 4GB memory, too much data was being gathered in the tables to answer the query. In more detail, we used the queries 1, 2, 3, 4, 5, 7, 10 from the LUBM benchmark. The results are shown for some representatives in Fig. 3. Basically, for queries 1, 3, 4, and 10, no real difference between EL and QL exists and the response time is strictly below 1s. For query 7, there exists a slight difference in favor of EL with 2.5s vs. 1s for LUBM10, whereas for queries 2 and 5 the difference increases. In all cases, the response time grows linearly w.r.t. the increasing size of LUBM, and we can conclude that on average querying in QL is slightly slower. Here, EL compensates for the longer preprocessing, and this effect becomes more visible, the more complex the query is and the more data needs to be gathered to answer it. Intuitively, this can be explained by looking at a simple example with two axioms $A \sqsubseteq \exists R$ and $\exists R \sqsubseteq B$. For EL, classification, yields $A \sqsubseteq B$ and only one axiom is translated and only one derivation step is required in XSB to obtain, say $B(a)$ from $A(a)$. For QL, both axioms are translated directly without classification, using DR , but now, two derivation steps would be required in XSB to obtain $B(a)$ from $A(a)$. It thus seems that deciding which of the two forms of translations performs better depends on the kind (and number) of queries we pose.

Finally, with the aim of also testing a more expressive

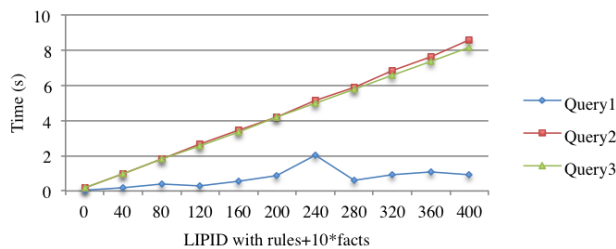


Figure 4: Query time for LIPID

OWL 2 QL ontology, we used the LIPID ontology,⁹ which has, besides 749 subclass axioms, 1,486 class disjointness axioms and 20 inverse object properties in combination with non-monotonic rules. The latter were created by means of the rule generator already used in [Ivanov *et al.*, 2013] with a ratio 1:10 between rules and facts, also introducing some new predicates not present in the ontology itself. We performed the preprocessing step and observed only slight effects due to the increasing amount of rules. The time for processing the ontology was naturally stable for all steps, and overall processing time was between 2 and 3s. Notably, the considerable amount of negative inclusions had no significant impact on time, e.g., when constructing the graph. Then, we posed three simple queries (Query1–3), namely *Acyl_Ester_Chain(X)*, *Lipid(X)*, and *Entity(X)* to the resulting rule sets in XSB. The results are shown in Fig. 4. As we can see, the response time is still very reasonable, from clearly below 1s to up to 8s. Still, the results in our opinion already show the effect of the arbitrary rules that tend to introduce links between predicates that increase the search space. This can be noted in particular for Query1, where in one case a smaller set of rules results in a higher response time, simply because no generated rule set is a subset of another. We note that performance tests of querying (non-monotonic) rules and ontologies would considerably benefit from real datasets but to the best of our knowledge currently none are available.

6 Conclusions

We have extended NoHR, the Protégé plug-in that allows to query non-monotonic rules and ontologies in OWL 2 EL, to also admit ontologies in OWL 2 QL. While the principal architecture of the tool remains the same, the crucial module that translates the ontology into rules with the help of a classifier simply cannot be re-used, which is why we introduced a novel direct translation for OWL 2 QL ontologies to cover this profile. We have implemented this translation and discussed optimizations. The evaluation shows that it maintains all positive evaluation results of the OWL 2 EL version [Ivanov *et al.*, 2013], and is even faster during pre-processing, as no classification is necessary, in exchange for an on average slightly longer response time during querying.

Besides the OWL 2 EL profile supported by NoHR, and compared to in Sect. 5, also [Gomes *et al.*, 2010; Knorr and

Alferes, 2011] both build on the well-founded MKNF semantics [Knorr *et al.*, 2011]. While [Gomes *et al.*, 2010] uses the non-standard CDF framework integrated in XSB, which complicates compatibility to standard OWL tools based on the OWL API, [Knorr and Alferes, 2011] presents an OWL 2 QL oracle based on common rewritings in the underlying DL *DL-Lite_R* [Artale *et al.*, 2009], but would require constant interaction between a rule reasoner and a DL reasoner, which is why we believe it to be less efficient than our approach.

Two related tools are DReW [Xiao *et al.*, 2013] and HD Rules [Drabent *et al.*, 2007], although based on different underlying formalisms to combine ontologies and rules (c.f. [Eiter *et al.*, 2008; Motik and Rosati, 2010] for a comparison), which, again, considerably complicates comparison.

Future work includes the extension to OWL 2 RL, but developing an alternative for OWL 2 QL using the classifier integrated in ontop [Kontchakov *et al.*, 2014] once its OWL API becomes available, or even the general reasoner Konclude [Steigmiller *et al.*, 2014], could shed more light on whether classification or direct translation fares better for proper OWL 2 QL ontologies. The efficiency of the latter reasoner also motivates looking into non-polynomial DLs, with possible influences from recent work on rewriting disjunctive datalog programs [Kaminski *et al.*, 2014]. Finally, we may extend NoHR for OWL 2 QL (and EL) to the paraconsistent semantics [Kaminski *et al.*, 2015] that would provide true support to the already occasionally observed paraconsistent behavior, or alternatively, to either generalizations of hybrid KBs [Gonçalves and Alferes, 2010; Knorr *et al.*, 2014; 2012; Knorr, 2015] or dynamics in hybrid KBs [Slota *et al.*, 2011; Slota and Leite, 2012], or even both [Gonçalves *et al.*, 2014].

References

- [Alberti *et al.*, 2012] M. Alberti, M. Knorr, A. S. Gomes, J. Leite, R. Gonçalves, and M. Slota. Normative systems require hybrid knowledge bases. In *Procs. of AAMAS. IFAAMAS*, 2012.
- [Alferes *et al.*, 2013] J. J. Alferes, M. Knorr, and T. Swift. Query-driven procedures for hybrid MKNF knowledge bases. *ACM Trans. Comput. Log.*, 14(2):1–43, 2013.
- [Artale *et al.*, 2009] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev. The *DL-Lite* family and relations. *J. Artif. Intell. Res. (JAIR)*, 36:1–69, 2009.
- [Baader *et al.*, 2005] F. Baader, S. Brandt, and C. Lutz. Pushing the \mathcal{EL} envelope. In *Procs. of IJCAI*. Professional Book Center, 2005.
- [Baader *et al.*, 2010] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2010.
- [Boley and Kifer, 2013] H. Boley and M. Kifer, editors. *RIF Overview*. W3C Recommendation, 5 February, 2013. Available at <http://www.w3.org/TR/rif-overview/>.
- [Calvanese *et al.*, 2007] D. Calvanese, G. de Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *Journal of Automated Reasoning*, 39(3):385–429, 2007.
- [Calvanese *et al.*, 2011] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, M. Rodriguez-Muro, R. Rosati, M. Ruzzi, and D. F. Savo. The MASTRO system for ontology-based data access. *Semantic Web*, 2(1):43–53, 2011.

⁹<http://bioonto.dcs.aber.ac.uk/ql-ont/>

- [Chen and Warren, 1996] W. Chen and D. S. Warren. Tabled Evaluation with Delaying for General Logic Programs. *Journal of the ACM*, 43(1):20–74, 1996.
- [Drabent *et al.*, 2007] W. Drabent, J. Henriksson, and J. Maluszynski. Hd-rules: A hybrid system interfacing prolog with dl-reasoners. In *Procs. of ALPSWS*. CEUR-WS.org, 2007.
- [Eiter *et al.*, 2008] T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining answer set programming with description logics for the semantic web. *Artif. Intell.*, 172(12-13):1495–1539, 2008.
- [Gelder *et al.*, 1991] A. Van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *J. ACM*, 38(3):620–650, 1991.
- [Gomes *et al.*, 2010] A. Sofia Gomes, J. J. Alferes, and T. Swift. Implementing query answering for hybrid MKNF knowledge bases. In *Procs. of PADL*. Springer, 2010.
- [Gonçalves *et al.*, 2014] R. Gonçalves, M. Knorr, and J. Leite. Evolving multi-context systems. In *Procs. of ECAI*. IOS Press, 2014.
- [Gonçalves and Alferes, 2010] R. Gonçalves and J. Alferes. Parametrized logic programming. In *Procs. of JELIA*. Springer, 2010.
- [Guo *et al.*, 2005] Y. Guo, Z. Pan, and J. Heflin. LUBM: A benchmark for OWL knowledge base systems. *J. Web Sem.*, 3(2-3):158–182, 2005.
- [Ivanov *et al.*, 2013] V. Ivanov, M. Knorr, and J. Leite. A query tool for \mathcal{EL} with non-monotonic rules. In *Procs. of ISWC*. Springer, 2013.
- [Kaminski *et al.*, 2014] M. Kaminski, Y. Nenov, and B. Cuenca Grau. Datalog rewritability of disjunctive datalog programs and its applications to ontology reasoning. In *Procs. of AAAI*. AAAI Press, 2014.
- [Kaminski *et al.*, 2015] T. Kaminski, M. Knorr, and J. Leite. Efficient paraconsistent reasoning with ontologies and rules. In *Procs. of IJCAI*, 2015.
- [Kazakov *et al.*, 2013] Y. Kazakov, M. Krötzsch, and F. Simančík. The incredible ELK: From polynomial procedures to efficient reasoning with \mathcal{EL} ontologies. *Journal of Automated Reasoning*, 53:1–61, 2013.
- [Knorr and Alferes, 2011] M. Knorr and J. J. Alferes. Querying OWL 2 QL and non-monotonic rules. In *Procs. of ISWC*. Springer, 2011.
- [Knorr *et al.*, 2011] M. Knorr, J. J. Alferes, and P. Hitzler. Local closed world reasoning with description logics under the well-founded semantics. *Artif. Intell.*, 175(9–10):1528–1554, 2011.
- [Knorr *et al.*, 2012] M. Knorr, P. Hitzler, and F. Maier. Reconciling OWL and non-monotonic rules for the semantic web. In *Procs. of ECAI*. IOS Press, 2012.
- [Knorr *et al.*, 2014] M. Knorr, M. Slota, J. Leite, and M. Homola. What if no hybrid reasoner is available? hybrid MKNF in multi-context systems. *J. Log. Comput.*, 24(6):1279–1311, 2014.
- [Knorr, 2015] Matthias Knorr. Nonmonotonic nominal schemas revisited. In *Procs. of DL*. CEUR-WS.org, 2015.
- [Kontchakov *et al.*, 2011] R. Kontchakov, C. Lutz, D. Toman, F. Wolter, and M. Zakharyashev. The combined approach to ontology-based data access. In *Procs. of IJCAI*. IJCAI/AAAI, 2011.
- [Kontchakov *et al.*, 2014] R. Kontchakov, M. Rezk, M. Rodriguez-Muro, G. Xiao, and M. Zakharyashev. Answering SPARQL queries over databases under OWL 2 QL entailment regime. In *Procs. of ISWC*. Springer, 2014.
- [Lembo *et al.*, 2013] D. Lembo, V. Santarelli, and D. F. Savo. Graph-based ontology classification in OWL 2 QL. In *Procs. of ESWC*. Springer, 2013.
- [Lifschitz, 1991] V. Lifschitz. Nonmonotonic databases and epistemic queries. In *Procs. of IJCAI*. Morgan Kaufmann, 1991.
- [Motik and Rosati, 2010] B. Motik and R. Rosati. Reconciling description logics and rules. *J. ACM*, 57(5):93–154, 2010.
- [Motik *et al.*, 2013] B. Motik, B. Cuenca Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz, editors. *OWL 2 Web Ontology Language: Profiles*. W3C Recommendation, 5 February, 2013. Available at <http://www.w3.org/TR/owl2-profiles/>.
- [Patel *et al.*, 2007] C. Patel, J. J. Cimino, J. Dolby, A. Fokoue, A. Kalyanpur, A. Kershenbaum, L. Ma, E. Schonberg, and K. Srinivas. Matching patient records to clinical trials using ontologies. In *Procs. of ISWC/ASWC*. Springer, 2007.
- [Savo *et al.*, 2010] D. F. Savo, D. Lembo, M. Lenzerini, A. Poggi, M. Rodriguez-Muro, V. Romagnoli, M. Ruzzi, and G. Stella. Mastro at work: Experiences on ontology-based data access. In *Procs. of DL*. CEUR-WS.org, 2010.
- [Slota and Leite, 2012] M. Slota and J. Leite. A unifying perspective on knowledge updates. In *Procs. of JELIA*. Springer, 2012.
- [Slota *et al.*, 2011] M. Slota, J. Leite, and T. Swift. Splitting and updating hybrid knowledge bases. *TPLP*, 11(4-5):801–819, 2011.
- [Steigmiller *et al.*, 2014] A. Steigmiller, T. Liebig, and B. Glimm. Konclude: System description. *J. Web Sem.*, 27:78–85, 2014.
- [Xiao *et al.*, 2013] G. Xiao, T. Eiter, and S. Heymans. The DReW system for nonmonotonic dl-programs. In *Procs. of SWWS*. Springer, 2013.