

Ontology Module Extraction via Datalog Reasoning

Ana Armas Romero and Mark Kaminski and Bernardo Cuenca Grau and Ian Horrocks

Department of Computer Science, University of Oxford, UK

Abstract

Module extraction—the task of computing a (preferably small) fragment $\mathcal{M}$ of an ontology $\mathcal{T}$ that preserves entailments over a signature Σ —has found many applications in recent years. Extracting modules of minimal size is, however, computationally hard, and often algorithmically infeasible. Thus, practical techniques are based on approximations, where $\mathcal{M}$ provably captures the relevant entailments, but is not guaranteed to be minimal. Existing approximations, however, ensure that $\mathcal{M}$ preserves all second-order entailments of $\mathcal{T}$ w.r.t. Σ , which is stronger than is required in many applications, and may lead to large modules in practice. In this paper we propose a novel approach in which module extraction is reduced to a reasoning problem in datalog. Our approach not only generalises existing approximations in an elegant way, but it can also be tailored to preserve only specific kinds of entailments, which allows us to extract significantly smaller modules. An evaluation on widely-used ontologies has shown very encouraging results.

1 Introduction

Module extraction is the task of computing, given an ontology $\mathcal{T}$ and a signature of interest Σ , a (preferably small) subset $\mathcal{M}$ of $\mathcal{T}$ (a module) that preserves all relevant entailments in $\mathcal{T}$ over the set of symbols Σ . Such an $\mathcal{M}$ is indistinguishable from $\mathcal{T}$ w.r.t. Σ , and $\mathcal{T}$ can be safely replaced with $\mathcal{M}$ in applications of $\mathcal{T}$ that use only the symbols in Σ .

Module extraction has received a great deal of attention in recent years (Stuckenschmidt, Parent, and Spaccapietra 2009; Cuenca Grau et al. 2008; Seidenberg and Recitor 2006; Kontchakov, Wolter, and Zakharyashev 2010; Gatens, Konev, and Wolter 2014; Del Vescovo et al. 2011; Nortje, Britz, and Meyer 2013), and modules have found a wide range of applications, including ontology reuse (Cuenca Grau et al. 2008; Jiménez-Ruiz et al. 2008), matching (Jiménez-Ruiz and Cuenca Grau 2011), debugging (Suntisivaraporn et al. 2008; Ludwig 2014) and classification (Armas Romero, Cuenca Grau, and Horrocks 2012; Tsarkov and Palmisano 2012; Cuenca Grau et al. 2010).

The preservation of relevant entailments is formalised via *inseparability relations*. The strongest notion is *model inseparability*, which requires that it must be possible to turn any model of $\mathcal{M}$ into a model of $\mathcal{T}$ by (re-)interpreting only the symbols outside Σ ; such an $\mathcal{M}$ preserves all second-order

entailments of $\mathcal{T}$ w.r.t. Σ (Konev et al. 2013). A weaker and more flexible notion is *deductive inseparability*, which requires only that $\mathcal{T}$ and $\mathcal{M}$ entail the same Σ -formulas *in a given query language*. Unfortunately, the decision problems associated with module extraction are invariably of high complexity, and often undecidable. For model inseparability, checking whether $\mathcal{M}$ is a Σ -module in $\mathcal{T}$ is undecidable even if $\mathcal{T}$ is restricted to be in the description logic (DL) $\mathcal{EL}$, for which standard reasoning is tractable. For deductive inseparability, the problem is typically decidable for lightweight DLs and “reasonable” query languages, albeit of high worst-case complexity; e.g., the problem is already EXPTIME-hard for $\mathcal{EL}$ if we consider concept inclusions as the query language (Lutz and Wolter 2010). Practical algorithms that ensure minimality of the extracted modules are known only for acyclic $\mathcal{ELI}$ (Konev et al. 2013) and DL-Lite (Kontchakov, Wolter, and Zakharyashev 2010).

Practical module extraction techniques are typically based on sound approximations: they ensure that the extracted fragment $\mathcal{M}$ is a module (i.e., inseparable from $\mathcal{T}$ w.r.t. Σ), but they give no minimality guarantee. The most popular such techniques are based on a family of polynomially checkable conditions called syntactic locality (Cuenca Grau et al. 2007; 2008; Sattler, Schneider, and Zakharyashev 2009); in particular, $\perp$ -locality and $\top\perp^*$ -locality. Each locality-based module $\mathcal{M}$ enjoys a number of desirable properties for applications: (i) it is model inseparable from $\mathcal{T}$; (ii) it is *depleting*, in the sense that $\mathcal{T} \setminus \mathcal{M}$ is inseparable from the empty ontology w.r.t. Σ ; (iii) it contains all justifications (a.k.a. explanations) in $\mathcal{T}$ of every Σ -formula entailed by $\mathcal{T}$; and (iv) last but not least, it can be computed efficiently, even for very expressive ontology languages.

Locality-based techniques are easy to implement, and surprisingly effective in practice. Their main drawback is that the extracted modules can be rather large, which limits their usefulness in some applications (Del Vescovo et al. 2013). One way to address this issue is to develop techniques that more closely approximate minimal modules while still preserving properties (i)–(iii). Efforts in this direction have confirmed that locality-based modules can be far from optimal in practice (Gatens, Konev, and Wolter 2014); however, these techniques apply only to rather restricted ontology languages and utilise algorithms with high worst-case complexity.

Another approach to computing smaller modules is to weaken properties (i)–(iii), which are stronger than is required in many applications. In particular, model inseparability (property (i)) is a very strong condition, and deductive inseparability would usually suffice, with the query language determining which kinds of consequence are preserved; in modular classification, for example, only atomic concept inclusions need to be preserved. However, all practical module extraction techniques that are applicable to expressive ontology languages yield modules satisfying all three properties, and hence potentially much larger than they need to be.

In this paper, we propose a technique that reduces module extraction to a reasoning problem in datalog. The connection between module extraction and datalog was first observed in (Suntisrivaraporn 2008), where it was shown that locality $\perp$ -module extraction for $\mathcal{EL}$ ontologies could be reduced to propositional datalog reasoning. Our approach takes this connection much farther by generalising both locality-based and reachability-based (Nortje, Britz, and Meyer 2013) modules for expressive ontology languages in an elegant way. A key distinguishing feature of our technique is that it can extract deductively inseparable modules, with the query language tailored to the requirements of the application at hand, which allows us to relax Property (i) and extract significantly smaller modules. In all cases our modules preserve the nice features of locality: they are widely applicable (even beyond DLs), they can be efficiently computed, they are depleting (Property (ii)) and they preserve all justifications of relevant entailments (Property (iii)).

We have implemented our approach using the RDFS engine (Motik et al. 2014). Our proof of concept evaluation shows that module size consistently decreases as we consider weaker inseparability relations, which could significantly improve the usefulness of modules in applications.

All our proofs are deferred to the appendix.

2 Preliminaries

Ontologies and Queries We use standard first-order logic and assume familiarity with description logics, ontology languages and theorem proving. A signature Σ is a set of predicates and $\text{Sig}(F)$ denotes the signature of a set of formulas F . It is assumed that the nullary falsehood predicate $\perp$ belongs to every Σ . To capture a wide range of KR languages, we formalise ontology axioms as *rules*: function-free sentences of the form $\forall \mathbf{x}. [\varphi(\mathbf{x}) \rightarrow \exists \mathbf{y}. [\bigvee_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})]]$, where φ, ψ_i are conjunctions of distinct atoms. Formula φ is the rule *body* and $\exists \mathbf{y}. [\bigvee_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})]$ is the *head*. Universal quantification is omitted for brevity. Rules are required to be safe (all variables in the head occur in the body) and we assume w.l.o.g. that $\top$ (resp. $\perp$) does not occur in rule heads (resp. in rule bodies). A TBox $\mathcal{T}$ is a finite set of rules; TBoxes mentioning equality ($\approx$) are extended with its standard axiomatisation. A fact γ is a function-free ground atom. An ABox $\mathcal{A}$ is a finite set of facts. A *positive existential query* (PEQ) is a formula $q(\mathbf{x}) = \exists \mathbf{y}. \varphi(\mathbf{x}, \mathbf{y})$, where φ is built from function-free atoms using only $\wedge$ and $\vee$.

Datalog A rule is *datalog* if its head has at most one atom and all variables are universally quantified. A *datalog pro-*

gram $\mathcal{P}$ is a set of datalog rules. Given $\mathcal{P}$ and an ABox $\mathcal{A}$, their *materialisation* is the set of facts entailed by $\mathcal{P} \cup \mathcal{A}$, which can be computed by means of forward-chaining. A fact γ is a consequence of a datalog rule $r = \bigwedge_{i=1}^n \gamma'_i \rightarrow \delta$ and facts $\gamma_1, \dots, \gamma_n$ if $\gamma = \delta\sigma$ with σ a most-general unifier (MGU) of γ_i, γ'_i for each $1 \leq i \leq n$. A (forward-chaining) *proof* of γ in $\mathcal{P} \cup \mathcal{A}$ is a pair $\rho = (T, \lambda)$ where T is a tree, λ is a mapping from nodes in T to facts, and from edges in T to rules in $\mathcal{P}$, such that for each node v the following holds: 1. $\lambda(v) = \gamma$ if v is the root of T ; 2. $\lambda(v) \in \mathcal{A}$ if v is a leaf; and 3. if v has children $w_1, \dots, w_n$ then each edge from v to w_i is labelled by the same rule r and $\lambda(v)$ is a consequence of r and $\lambda(w_1), \dots, \lambda(w_n)$. Forward-chaining is sound and complete: a fact γ is in the materialisation of $\mathcal{P} \cup \mathcal{A}$ iff it has a proof in $\mathcal{P} \cup \mathcal{A}$. Finally, the *support* of γ is the set of rules occurring in some proof of γ in $\mathcal{P} \cup \mathcal{A}$.

Inseparability Relations & Modules We next recapitulate the most common inseparability relations studied in the literature. We say that TBoxes $\mathcal{T}$ and $\mathcal{T}'$ are

- Σ -*model inseparable* ($\mathcal{T} \equiv_{\Sigma}^m \mathcal{T}'$), if for every model $\mathcal{I}$ of $\mathcal{T}$ (resp. of $\mathcal{T}'$) there exists a model $\mathcal{J}$ of $\mathcal{T}'$ (resp. of $\mathcal{T}$) with the same domain s.t. $A^{\mathcal{I}} = A^{\mathcal{J}}$ for each $A \in \Sigma$.
- Σ -*query inseparable* ($\mathcal{T} \equiv_{\Sigma}^q \mathcal{T}'$) if for every Boolean PEQ q and Σ -ABox $\mathcal{A}$ we have $\mathcal{T} \cup \mathcal{A} \models q$ iff $\mathcal{T}' \cup \mathcal{A} \models q$.
- Σ -*fact inseparable* ($\mathcal{T} \equiv_{\Sigma}^f \mathcal{T}'$) if for every fact γ and ABox $\mathcal{A}$ over Σ we have $\mathcal{T} \cup \mathcal{A} \models \gamma$ iff $\mathcal{T}' \cup \mathcal{A} \models \gamma$.
- Σ -*implication inseparable* ($\mathcal{T} \equiv_{\Sigma}^i \mathcal{T}'$) if for each φ of the form $A(\mathbf{x}) \rightarrow B(\mathbf{x})$ with $A, B \in \Sigma$, $\mathcal{T} \models \varphi$ iff $\mathcal{T}' \models \varphi$.

These relations are naturally ordered from strongest to weakest: $\equiv_{\Sigma}^m \subseteq \equiv_{\Sigma}^q \subseteq \equiv_{\Sigma}^f \subseteq \equiv_{\Sigma}^i$ for each non-trivial Σ .

Given an inseparability relation $\equiv$ for Σ , a subset $\mathcal{M} \subseteq \mathcal{T}$ is a $\equiv$ -*module* of $\mathcal{T}$ if $\mathcal{T} \equiv \mathcal{M}$. Furthermore, $\mathcal{M}$ is *minimal* if no $\mathcal{M}' \subsetneq \mathcal{M}$ is a $\equiv$ -module of $\mathcal{T}$.

3 Module Extraction via Datalog Reasoning

In this section, we present our approach to module extraction by reduction into a reasoning problem in datalog. Our approach builds on recent techniques that exploit datalog engines for ontology reasoning (Kontchakov et al. 2011; Stefanoni, Motik, and Horrocks 2013; Zhou et al. 2014). In what follows, we fix an arbitrary TBox $\mathcal{T}$ and signature $\Sigma \subseteq \text{Sig}(\mathcal{T})$. Unless otherwise stated, our definitions and theorems are parameterised by such $\mathcal{T}$ and Σ . We assume w.l.o.g. that rules in $\mathcal{T}$ do not share existentially quantified variables. For simplicity, we also assume that $\mathcal{T}$ contains no constants (all our results can be seamlessly extended).

3.1 Overview and Main Intuitions

Our overall strategy to extract a module $\mathcal{M}$ of $\mathcal{T}$ for an inseparability relation $\equiv_{\Sigma}^z$, with $z \in \{m, q, f, i\}$, can be summarised by the following steps:

1. Pick a substitution θ mapping all existentially quantified variables in $\mathcal{T}$ to constants, and transform $\mathcal{T}$ into a datalog program $\mathcal{P}$ by (i) Skolemising all rules in $\mathcal{T}$ using θ and (ii) turning disjunctions into conjunctions while splitting them into different rules, thus replacing each function-free

(r_1)	$A(x) \rightarrow \exists y_1.[R(x, y_1) \wedge B(y_1)]$	$A \sqsubseteq \exists R.B$
(r_2)	$A(x) \rightarrow \exists y_2.[R(x, y_2) \wedge C(y_2)]$	$A \sqsubseteq \exists R.C$
(r_3)	$B(x) \wedge C(x) \rightarrow D(x)$	$B \sqcap C \sqsubseteq D$
(r_4)	$D(x) \rightarrow \exists y_3.[S(x, y_3) \wedge E(y_3)]$	$D \sqsubseteq \exists S.E$
(r_5)	$D(x) \wedge S(x, y) \rightarrow F(y)$	$D \sqsubseteq \forall S.F$
(r_6)	$S(x, y) \wedge E(y) \wedge F(y) \rightarrow G(x)$	$\exists S.(E \sqcap F) \sqsubseteq G$
(r_7)	$G(x) \wedge H(x) \rightarrow \perp$	$G \sqcap H \sqsubseteq \perp$

Figure 1: Example TBox $\mathcal{T}^{ex}$ with DL translation

disjunctive rule of the form $\varphi(\mathbf{x}) \rightarrow \bigvee_{i=1}^n \psi_i(\mathbf{x})$ with datalog rules $\varphi(\mathbf{x}) \rightarrow \psi_1(\mathbf{x}), \dots, \varphi(\mathbf{x}) \rightarrow \psi_n(\mathbf{x})$.

2. Pick a Σ -ABox $\mathcal{A}_0$ and materialise $\mathcal{P} \cup \mathcal{A}_0$.
3. Pick a set $\mathcal{A}_r$ of “relevant facts” in the materialisation and compute the supporting rules in $\mathcal{P}$ for each such fact.
4. The module $\mathcal{M}$ consists of all rules in $\mathcal{T}$ that yield some supporting rule in $\mathcal{P}$. In this way, $\mathcal{M}$ is fully determined by the substitution θ and the ABoxes $\mathcal{A}_0$ and $\mathcal{A}_r$.

The main intuition behind our module extraction approach is that we can pick θ , $\mathcal{A}_0$ and $\mathcal{A}_r$ (and hence $\mathcal{M}$) such that each proof ρ of a Σ -consequence φ of $\mathcal{T}$ to be preserved can be embedded in a forward chaining proof ρ' in $\mathcal{P} \cup \mathcal{A}_0$ of a relevant fact in $\mathcal{A}_r$. Such an embedding satisfies the key property that, for each rule r involved in ρ , at least one corresponding datalog rule in $\mathcal{P}$ is involved in ρ' . In this way we ensure that $\mathcal{M}$ contains the necessary rules to entail φ . This approach, however, does not ensure minimality of $\mathcal{M}$: since $\mathcal{P}$ is a strengthening of $\mathcal{T}$ there may be proofs of a relevant fact in $\mathcal{P} \cup \mathcal{A}_0$ that do not correspond to a Σ -consequence of $\mathcal{T}$, which may lead to unnecessary rules in $\mathcal{M}$.

To illustrate how our strategy might work in practice, suppose that $\mathcal{T}$ is $\mathcal{T}^{ex}$ in Fig. 1, $\Sigma = \{B, C, D, G\}$, and that we want a module $\mathcal{M}$ that is Σ -implication inseparable from $\mathcal{T}^{ex}$. This is a simple case since $\varphi = D(x) \rightarrow G(x)$ is the only non-trivial Σ -implication entailed by $\mathcal{T}^{ex}$; thus, for $\mathcal{M}$ to be a module we only require that $\mathcal{M} \models \varphi$.

Proving $\mathcal{T}^{ex} \models \varphi$ amounts to proving $\mathcal{T}^{ex} \cup \{D(a)\} \models G(a)$ (with a a fresh constant). Figure 2(a) depicts a hyper-resolution tree ρ showing how $G(a)$ can be derived from the clauses corresponding to r_4 – r_6 and $D(a)$, with rule r_4 transformed into clauses

$$r'_4 = D(x) \rightarrow S(x, f(x_3)) \quad r''_4 = D(x) \rightarrow E(f(x_3))$$

Hence $\mathcal{M} = \{r_4$ – $r_6\}$ is a Σ -implication inseparable module of $\mathcal{T}^{ex}$, and as $G(a)$ cannot be derived from any subset of $\{r_4$ – $r_6\}$, $\mathcal{M}$ is also minimal.

In our approach, we pick $\mathcal{A}_0$ to contain the initial fact $D(a)$, $\mathcal{A}_r$ to contain the fact to be proved $G(a)$, and we make θ map variable y_3 in r_4 to a fresh constant c , in which case rule r_4 corresponds to the following datalog rules in $\mathcal{P}$:

$$D(x) \rightarrow S(x, c) \quad D(x) \rightarrow E(c)$$

Figure 2(b) depicts a forward chaining proof ρ' of $G(a)$ in $\mathcal{P} \cup \{D(a)\}$. As shown in the figure, ρ can be embedded in ρ' via θ by mapping functional terms over f to the fresh constant c . In this way, the rules involved in ρ are mapped to the datalog rules involved in ρ' via θ . Consequently, we will extract the (minimal) module $\mathcal{M} = \{r_4$ – $r_6\}$.

3.2 The Notion of Module Setting

The substitution θ and the ABoxes $\mathcal{A}_0$ and $\mathcal{A}_r$, which determine the extracted module, can be chosen in different ways to ensure the preservation of different kinds of Σ -consequences. The following notion of a module setting captures in a declarative way the main elements of our approach.

Definition 1. A *module setting* for $\mathcal{T}$ and Σ is a tuple $\chi = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ with θ a substitution from existentially quantified variables in $\mathcal{T}$ to constants, $\mathcal{A}_0$ a Σ -ABox, $\mathcal{A}_r$ a $\text{Sig}(\mathcal{T})$ -ABox, and s.t. no constant in χ occurs in $\mathcal{T}$.

The *program* of χ is the smallest datalog program $\mathcal{P}^\chi$ containing, for each $r = \varphi(\mathbf{x}) \rightarrow \exists \mathbf{y}.[\bigvee_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})]$ in $\mathcal{T}$, the rule $\varphi \rightarrow \perp$ if $n = 0$ and all rules $\varphi \rightarrow \gamma\theta$ for each $1 \leq i \leq n$ and each atom γ in ψ_i . The *support* of χ is the set of rules $r \in \mathcal{P}^\chi$ that support a fact from $\mathcal{A}_r$ in $\mathcal{P}^\chi \cup \mathcal{A}_0$. The *module* $\mathcal{M}^\chi$ of χ is the set of rules in $\mathcal{T}$ that have a corresponding datalog rule in the support of χ . $\diamond$

3.3 Modules for each Inseparability Relation

We next consider each inseparability relation $\equiv_\Sigma^z$, where $z \in \{m, q, f, i\}$, and formulate a specific setting χ_z which provably yields a $\equiv_\Sigma^z$ -module of $\mathcal{T}$.

Implication Inseparability The example in Section 3.1 suggests a natural setting $\chi_i = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ that guarantees implication inseparability. As in our example, we pick θ to be as “general” as possible by Skolemising each existentially quantified variable to a fresh constant. For A and B predicates of the same arity n , proving that $\mathcal{T}$ entails a Σ -implication $\varphi = A(x_1, \dots, x_n) \rightarrow B(x_1, \dots, x_n)$, amounts to showing that $\mathcal{T} \cup \{A(a_1, \dots, a_n)\} \models B(a_1, \dots, a_n)$ for fresh constants $a_1, \dots, a_n$. Thus, following the ideas of our example, we initialise $\mathcal{A}_0$ with a fact $A(c_A^1, \dots, c_A^n)$ for each n -ary predicate $A \in \Sigma$, and $\mathcal{A}_r$ with a fact $B(c_A^1, \dots, c_A^n)$ for each pair of n -ary predicates $\{B, A\} \subseteq \Sigma$ with $B \neq A$.

Definition 2. For each existentially quantified variable y_j in $\mathcal{T}$, let c_{y_j} be a fresh constant. Furthermore, for each $A \in \Sigma$ of arity n , let $c_A^1, \dots, c_A^n$ be also fresh constants. The setting $\chi_i = \langle \theta^i, \mathcal{A}_0^i, \mathcal{A}_r^i \rangle$ is defined as follows:

- $\theta^i = \{y_j \mapsto c_{y_j} \mid y_j \text{ existentially quantified in } \mathcal{T}\}$,
- $\mathcal{A}_0^i = \{A(c_A^1, \dots, c_A^n) \mid A \text{ } n\text{-ary predicate in } \Sigma\}$, and
- $\mathcal{A}_r^i = \{B(c_A^1, \dots, c_A^n) \mid A \neq B \text{ } n\text{-ary predicates in } \Sigma\}$. $\diamond$

The setting χ_i is reminiscent of the datalog encodings typically used to check whether a concept A is subsumed by concept B w.r.t. a “lightweight” ontology $\mathcal{T}$ (Krötzsch, Rudolph, and Hitzler 2008; Stefanoni, Motik, and Horrocks 2013). There, variables in rules are Skolemised as fresh constants to produce a datalog program $\mathcal{P}$ and it is then checked whether $\mathcal{P} \cup \{A(a)\} \models B(a)$.

Theorem 3. $\mathcal{M}^{\chi_i} \equiv_\Sigma^i \mathcal{T}$.

Fact Inseparability The setting χ_i in Def. 2 cannot be used to ensure fact inseparability. Consider again $\mathcal{T}^{ex}$ and $\Sigma = \{B, C, D, G\}$, for which $\mathcal{M}^{\chi_i} = \{r_4, r_5, r_6\}$. For $\mathcal{A} = \{B(a), C(a)\}$ we have $\mathcal{T}^{ex} \cup \mathcal{A} \models G(a)$ but $\mathcal{M}^{\chi_i} \cup \mathcal{A} \not\models G(a)$, and hence $\mathcal{M}^{\chi_i}$ is not fact inseparable from $\mathcal{T}^{ex}$.




Figure 2: Proofs of $G(a)$ from $D(a)$ in (a) $\mathcal{T}^{ex}$ and (b) the corresponding datalog program

More generally, $\mathcal{M}^{\chi_i}$ is only guaranteed to preserve Σ -fact entailments $\mathcal{T} \cup \mathcal{A} \models \gamma$ where $\mathcal{A}$ is a singleton. However, for a module to be fact inseparable from $\mathcal{T}$ it must preserve all Σ -facts when coupled with *any* Σ -ABox. We achieve this by choosing $\mathcal{A}_0$ to be the *critical ABox* for Σ , which consists of all facts that can be constructed using Σ and a single fresh constant (Marnette 2009). Every Σ -ABox can be homomorphically mapped into the critical Σ -ABox. In this way, we can show that all proofs of a Σ -fact in $\mathcal{T} \cup \mathcal{A}$ can be embedded in a proof of a relevant fact in $\mathcal{P}^{\chi} \cup \mathcal{A}_0$.

Definition 4. Let constants c_{y_i} be as in Def. 2, and let $*$ be a fresh constant. The setting $\chi_f = \langle \theta^f, \mathcal{A}_0^f, \mathcal{A}_r^f \rangle$ is defined as follows: (i) $\theta^f = \theta^i$, (ii) $\mathcal{A}_0^f = \{A(*, \dots, *) \mid A \in \Sigma\}$, and (iii) $\mathcal{A}_r^f = \mathcal{A}_0^f$. $\diamond$

The datalog programs for χ_i and χ_f coincide and hence the only difference between the two settings is in the definition of their corresponding ABoxes. In our example, both $\mathcal{A}_0^f$ and $\mathcal{A}_r^f$ contain facts $B(*), C(*), D(*),$ and $G(*).$ Clearly, $\mathcal{P}^{\chi_f} \cup \mathcal{A}_0 \models G(*)$ and the proof additionally involves rule r_3 . Thus $\mathcal{M}^{\chi_f} = \{r_3, r_4, r_5, r_6\}$.

Theorem 5. $\mathcal{M}^{\chi_f} \equiv_{\Sigma}^f \mathcal{T}$.

Query Inseparability Positive existential queries constitute a much richer query language than facts as they allow for existentially quantified variables. Thus, the query inseparability requirement invariably leads to larger modules.

For instance, let $\mathcal{T} = \mathcal{T}^{ex}$ and $\Sigma = \{A, B\}$. Given the Σ -ABox $\mathcal{A} = \{A(a)\}$ and Σ -query $q = \exists y. B(y)$ we have that $\mathcal{T}^{ex} \cup \mathcal{A} \models q$ (due to rule r_1). The module $\mathcal{M}^{\chi_f}$ is, however, empty. Indeed, the materialisation of $\mathcal{P}^{\chi_f} \cup \{A(*)\}$ consists of the additional facts $R(*, c_{y_1})$ and $B(c_{y_1})$ and hence it does not contain any relevant fact mentioning only $*$. Thus, $\mathcal{M}^{\chi_f} \cup \mathcal{A} \not\models q$ and $\mathcal{M}^{\chi_f}$ is not query inseparable from $\mathcal{T}^{ex}$.

Our example suggests that, although the critical ABox is constrained enough to embed every Σ -ABox, we may need to consider additional relevant facts to capture all proofs of a Σ -query. In particular, rule r_1 implies that B contains an instance whenever A does: a dependency that is then checked by q . This can be captured by considering fact $B(c_{y_1})$ as relevant, in which case rule r_1 would be in the module.

More generally, we consider a module setting χ that differs from χ_f only in that all Σ -facts (and not just those over $*$) are considered as relevant.

Definition 6. Let constants c_{y_i} and $*$ be as in Def. 4. The setting $\chi_q = \langle \theta^q, \mathcal{A}_0^q, \mathcal{A}_r^q \rangle$ is as follows: (i) $\theta^q = \theta^f$, (ii) $\mathcal{A}_0^q = \mathcal{A}_0^f$, and (iii) $\mathcal{A}_r^q$ consists of all Σ -facts $A(a_1, \dots, a_n)$ with each a_j either a constant c_{y_i} or $*$. $\diamond$

Correctness is established by the following theorem:

Theorem 7. $\mathcal{M}^{\chi_q} \equiv_{\Sigma}^q \mathcal{T}$.

Model Inseparability The modules generated by χ_q may not be model inseparable from $\mathcal{T}$. To see this, let $\mathcal{T} = \mathcal{T}^{ex}$ and $\Sigma = \{A, D, R\}$, in which case $\mathcal{M}^{\chi_q} = \{r_1, r_2\}$. The interpretation $\mathcal{I}$ where $\Delta^{\mathcal{I}} = \{a, b\}$, $A^{\mathcal{I}} = \{a\}$, $B^{\mathcal{I}} = C^{\mathcal{I}} = \{b\}$, $D^{\mathcal{I}} = \emptyset$ and $R^{\mathcal{I}} = \{(a, b)\}$ is a model of $\mathcal{M}^{\chi_q}$. This interpretation, however, cannot be extended to a model of r_3 (and hence of $\mathcal{T}$) without reinterpreting A, R or D .

The main insight behind locality and reachability modules is to ensure that each model of the module can be extended to a model of $\mathcal{T}$ in a uniform way. Specifically, each model of a $\top \perp^*$ -locality or $\top \perp^*$ -reachability module can be extended to a model of $\mathcal{T}$ by interpreting all other predicates A as either $\emptyset$ or $(\Delta^{\mathcal{I}})^n$ with n the arity of A . Thus, $\mathcal{M} = \{r_1, r_2, r_3\}$ is a $\equiv_{\Sigma}^m$ -module of $\mathcal{T}^{ex}$ since all its models can be extended by interpreting E, F and G as the domain, H as empty, and S as the Cartesian product of the domain. We can capture this idea in our framework by means of the following setting.

Definition 8. The setting $\chi_m = \langle \theta^m, \mathcal{A}_0^m, \mathcal{A}_r^m \rangle$ is as follows: θ^m maps each existentially quantified variable to the fresh constant $*$ and $\mathcal{A}_0^m = \mathcal{A}_r^m = \mathcal{A}_0^f$. $\diamond$

In our example, $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ entails the relevant facts $A(*), R(*, *)$ and $D(*)$, and hence $\mathcal{M}^{\chi_m} = \{r_1, r_2, r_3\}$.

To show that $\mathcal{M}^{\chi_m}$ is a $\equiv_{\Sigma}^m$ -module we prove that all models $\mathcal{I}$ of $\mathcal{M}^{\chi_m}$ can be extended to a model of $\mathcal{T}$ as follows: (i) predicates not occurring in the materialisation of $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ are interpreted as empty; (ii) predicates in the support of χ_m (and hence occurring in $\mathcal{M}^{\chi_m}$) are interpreted as in $\mathcal{I}$; and (iii) all other predicates A are interpreted as $(\Delta^{\mathcal{I}})^n$ with n the arity of A .

Theorem 9. $\mathcal{M}^{\chi_m} \equiv_{\Sigma}^m \mathcal{T}$.

3.4 Modules for Ontology Classification

Module extraction has been exploited for optimising ontology classification (Armas Romero, Cuenca Grau, and Horrocks 2012; Tsarkov and Palmisano 2012; Cuenca Grau et al. 2010). In this case, it is not only required that modules

are implication inseparable from $\mathcal{T}$, but also that they preserve all implications $A(\mathbf{x}) \rightarrow B(\mathbf{x})$ with $A \in \Sigma$ but $B \notin \Sigma$. This requirement can be captured as given next.

Definition 10. TBoxes $\mathcal{T}$ and $\mathcal{T}'$ are Σ -classification inseparable ($\mathcal{T} \equiv_{\Sigma}^c \mathcal{T}'$) if for each φ of the form $A(\mathbf{x}) \rightarrow B(\mathbf{x})$ with $A \in \Sigma$, and $B \in \text{Sig}(\mathcal{T} \cup \mathcal{T}')$ we have $\mathcal{T} \models \varphi$ iff $\mathcal{T}' \models \varphi$. $\diamond$

Classification inseparability is a stronger requirement than implication inseparability. For $\mathcal{T} = \{A(x) \rightarrow B(x)\}$ and $\Sigma = \{A\}$, $\mathcal{M} = \emptyset$ is implication inseparable from $\mathcal{T}$, whereas classification inseparability requires that $\mathcal{M} = \mathcal{T}$.

Modular reasoners such as MORE and Chainsaw rely on locality $\perp$ -modules, which satisfy this requirement. Each model of a $\perp$ -module $\mathcal{M}$ can be extended to a model of $\mathcal{T}$ by interpreting all additional predicates as empty, which is not possible if $A \in \Sigma$ and $\mathcal{T}$ entails $A(x) \rightarrow B(x)$ but $\mathcal{M}$ does not. We can cast $\perp$ -modules in our framework with the following setting, which extends χ_m in Def. 8 by also considering as relevant facts involving predicates not in Σ .

Definition 11. The setting $\chi_b = \langle \theta^b, \mathcal{A}_0^b, \mathcal{A}_r^b \rangle$ is as follows: $\theta^b = \theta^m$, $\mathcal{A}_0^b = \mathcal{A}_0^m$, and $\mathcal{A}_r$ consists of all facts $A(*, \dots, *)$ where $A \in \text{Sig}(\mathcal{T})$. $\diamond$

The use of $\perp$ -modules is, however, stricter than is needed for ontology classification. For instance, if we consider $\mathcal{T} = \mathcal{T}^{ex}$ and $\Sigma = \{A\}$ we have that $\mathcal{M}^{\chi_b}$ contains all rules r_1 – r_6 , but since A does not have any subsumers in $\mathcal{T}^{ex}$ the empty TBox is already classification inseparable from $\mathcal{T}^{ex}$.

The following module setting extends χ_i in Def. 2 to ensure classification inseparability. As in the case of χ_b in Def. 11 the only required modification is to also consider as relevant facts involving predicates outside Σ .

Definition 12. Setting $\chi_c = \langle \theta^c, \mathcal{A}_0^c, \mathcal{A}_r^c \rangle$ is as follows: $\theta^c = \theta^i$, $\mathcal{A}_0^c = \mathcal{A}_0^i$, and $\mathcal{A}_r^c$ consists of all facts $B(c_A^1, \dots, c_A^n)$ s.t. $A \neq B$ are n -ary predicates, $A \in \Sigma$ and $B \in \text{Sig}(\mathcal{T})$. $\diamond$

Indeed, if we consider again $\mathcal{T} = \mathcal{T}^{ex}$ and $\Sigma = \{A\}$, the module for χ_c is empty, as desired.

Theorem 13. $\mathcal{M}^{\chi_c} \equiv_{\Sigma}^c \mathcal{T}$.

3.5 Additional Properties of Modules

Although the essential property of a module $\mathcal{M}$ is that it captures all relevant Σ -consequences of $\mathcal{T}$, in some applications it is desirable that modules satisfy additional requirements.

In ontology reuse scenarios, it is sometimes desirable that a module $\mathcal{M}$ does not “leave any relevant information behind”, in the sense that $\mathcal{T} \setminus \mathcal{M}$ does not entail any relevant Σ -consequence—a property referred to as *depleting-ness* (Kontchakov, Wolter, and Zakharyashev 2010).

Definition 14. Let $\equiv_{\Sigma}^z$ be an inseparability relation. A $\equiv_{\Sigma}^z$ -module $\mathcal{M}$ of $\mathcal{T}$ is *depleting* if $\mathcal{T} \setminus \mathcal{M} \equiv_{\Sigma}^z \emptyset$. $\diamond$

Note that not all modules are depleting: for some relevant Σ -entailment φ it may be that $\mathcal{M} \models \varphi$ (as required by the definition of module), but also that $(\mathcal{T} \setminus \mathcal{M}) \models \varphi$, in which case $\mathcal{M}$ is not depleting. The following theorem establishes that all modules defined in Section 3.3 are depleting.

Theorem 15. $\mathcal{M}^{\chi_z}$ is depleting for each $z \in \{m, q, f, i, c\}$.

Another common application of modules is to optimise the computation of justifications: minimal subsets of a TBox that are sufficient to entail a given formula (Kalyanpur et al. 2007; Suntisrivaraporn et al. 2008).

Definition 16. Let $\mathcal{T} \models \varphi$. A *justification* for φ in $\mathcal{T}$ is a minimal subset $\mathcal{T}' \subseteq \mathcal{T}$ such that $\mathcal{T}' \models \varphi$. $\diamond$

Justifications are displayed in ontology development platforms as explanations of why an entailment holds, and tools typically compute all of them. Extracting justifications is a computationally intensive task, and locality-based modules have been used to reduce the size of the problem: if $\mathcal{T}'$ is a justification of φ in $\mathcal{T}$, then $\mathcal{T}'$ is contained in a locality module of $\mathcal{T}$ for $\Sigma = \text{Sig}(\varphi)$. Our modules are also justification-preserving, and we can adjust our modules depending on what kind of first-order sentence φ is.

Theorem 17. Let $\mathcal{T}'$ be a justification for a first-order sentence φ in $\mathcal{T}$ and let $\text{Sig}(\varphi) \subseteq \Sigma$. Then, $\mathcal{T}' \subseteq \mathcal{M}^{\chi_m}$. Additionally, the following properties hold: (i) if φ is a rule, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_q}$; (ii) if φ is datalog, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_f}$; and (iii) if φ is of the form $A(\mathbf{x}) \rightarrow B(\mathbf{x})$, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_i}$; finally, if φ satisfies $A \in \Sigma, B \in \text{Sig}(\mathcal{T})$, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_c}$.

3.6 Complexity of Module Extraction

We conclude this section by showing that our modules can be efficiently computed in most practically relevant cases.

Theorem 18. Let m be a non-negative integer and L a class of TBoxes s.t. each rule in a TBox from L has at most m distinct universally quantified variables. The following problem is tractable: given $z \in \{q, f, i, c\}$, $\mathcal{T} \in L$, and $r \in \mathcal{T}$, decide whether $r \in \mathcal{M}^{\chi_z}$. The problem is solvable in polynomial time for arbitrary classes L of TBoxes if $z = m$.

We now provide a proof sketch for this result. Checking whether a datalog program $\mathcal{P}$ and an ABox $\mathcal{A}$ entail a fact is feasible in $\mathcal{O}(|\mathcal{P}| \cdot n^v)$, with n the number of constants in $\mathcal{P} \cup \mathcal{A}$ and v the maximum number of variables in a rule from $\mathcal{P}$ (Dantsin et al. 2001). Thus, although datalog reasoning is exponential in the size of v (and hence of $\mathcal{P}$), it is tractable if v is bounded by a constant.

Given arbitrary $\mathcal{T}$ and Σ , and for $z \in \{m, q, f, i, c\}$, the datalog program $\mathcal{P}^{\chi_z}$ can be computed in linear time in the size of $|\mathcal{T}|$. The number of constants n in χ_z (and hence in $\mathcal{P}^{\chi_z} \cup \mathcal{A}_0^z$) is linearly bounded in $|\mathcal{T}|$, whereas the maximum number of variables v coincides with the maximum number of universally quantified variables in a rule from $\mathcal{T}$. As shown in (Zhou et al. 2014), computing the support of a fact in a datalog program is no harder than fact entailment, and thus module extraction in our approach is feasible in $\mathcal{O}(|\mathcal{T}| \cdot n^v)$, and thus tractable for ontology languages where rules have a bounded number of variables (as is the case for most DLs). Finally, if $z = m$ the setting χ_m involves a single constant $*$ and module extraction boils down to reasoning in propositional datalog (a tractable problem regardless of $\mathcal{T}$).

3.7 Module Containment and Optimality

Intuitively, the more expressive the language for which preservation of consequences is required the larger modules need to be. The following proposition shows that our modules are consistent with this intuition.

Proposition 19. $\mathcal{M}^{\chi_i} \subseteq \mathcal{M}^{\chi_f} \subseteq \mathcal{M}^{\chi_q} \subseteq \mathcal{M}^{\chi_m} \subseteq \mathcal{M}^{\chi_b}$ and $\mathcal{M}^{\chi_i} \subseteq \mathcal{M}^{\chi_c} \subseteq \mathcal{M}^{\chi_b}$

As already discussed, these containment relations are strict for many $\mathcal{T}$ and Σ .

We conclude this section by discussing whether each χ_z with $z \in \{q, f, i, c\}$ is optimal for its inseparability relation in the sense that there is no setting that produces smaller modules. To make optimality statements precise we need to consider families of module settings, that is, functions that assign a module setting for each pair of $\mathcal{T}$ and Σ .

Definition 20. A *setting family* is a function Ψ that maps a TBox $\mathcal{T}$ and signature Σ to a module setting for $\mathcal{T}$ and Σ . We say that Ψ is *uniform* if for every Σ and pair of TBoxes $\mathcal{T}, \mathcal{T}'$ with the same number of existentially quantified variables $\Psi(\mathcal{T}, \Sigma) = \Psi(\mathcal{T}', \Sigma)$. Let $z \in \{i, f, q, c\}$; then, Ψ is *z-admissible* if, for each $\mathcal{T}$ and Σ , $\mathcal{M}^{\Psi(\mathcal{T}, \Sigma)}$ is a $\equiv_{\Sigma}^z$ -module of $\mathcal{T}$. Finally, Ψ is *z-optimal* if $\mathcal{M}^{\Psi(\mathcal{T}, \Sigma)} \subseteq \mathcal{M}^{\Psi'(\mathcal{T}, \Sigma)}$ for every $\mathcal{T}, \Sigma$ and every uniform Ψ' that is *z-admissible*. $\diamond$

Uniformity ensures that settings do not depend on the specific shape of rules in $\mathcal{T}$, but rather only on Σ and the number of existentially quantified variables in $\mathcal{T}$. In turn, admissibility ensures that each setting yields a module. The (uniform and admissible) family Ψ^z corresponding to each setting χ^z in Sections 3.3 and 3.4 is defined in the obvious way: for each $\mathcal{T}$ and Σ , $\Psi^z(\mathcal{T}, \Sigma)$ is the setting χ^z for $\mathcal{T}$ and Σ .

The next theorem shows that Ψ^z is optimal for implication and classification inseparability.

Theorem 21. Ψ^z is *z-optimal* for $z \in \{i, c\}$.

In contrast, Ψ^q and Ψ^f are not optimal. To see this, let $\mathcal{T} = \{A(x) \rightarrow B(x), B(x) \rightarrow A(x)\}$ and $\Sigma = \{A\}$. The empty TBox is fact inseparable from $\mathcal{T}$ since the only Σ -consequence of $\mathcal{T}$ is the tautology $A(x) \rightarrow A(x)$. However, $\mathcal{M}^{\chi_f} = \mathcal{T}$ since fact $A(a)$ is in $\mathcal{A}_r^f$ and its support is included in the module. We can provide a family of settings that distinguishes tautological from non-tautological inferences (see appendix); however, this family yields settings of exponential size in $|\mathcal{T}|$, which is undesirable in practice.

4 Proof of Concept Evaluation

We have implemented a prototype system for module extraction that uses RDFox for datalog materialisation (Motik et al. 2014). Additionally, the ontology reasoner PAGOdA (Zhou et al. 2014) provides functionality for computing the support of an entailed fact in datalog, which we have adapted for computing modules. We have evaluated our system on representative ontologies, including SNOMED (SCT), Fly Anatomy (FLY), the Gene Ontology (GO) and BioModels (BM).¹ SCT is expressed in the EL profile of OWL 2, whereas FLY, GO and BM require expressive DLs (Horn-SRI, SHIQ and SRIQ, respectively). We have normalised all ontologies to make axioms equivalent to rules.

We compared the size of our modules with the locality-based modules computed using the OWL API. We have followed the experimental methodology from (Del Vescovo et

¹The ontologies used in our tests are available for download at <http://www.cs.ox.ac.uk/isg/ontologies/UID/> under IDs 794 (FLY), 795 (SCT), 796 (GO) and 797 (BM).

	FLY		SCT		GO		BM	
rules	19,830		112,833		145,083		462,120	
	gen	rnd	gen	rnd	gen	rnd	gen	rnd
$\perp, \chi_b$	242	847	242	5,196	1,461	12,801	1,010	64,320
χ_c	112	446	230	3,500	309	3,990	285	14,273
$\top\perp^*$	219	796	233	5,182	1,437	12,747	963	62,897
χ_m	215	789	233	5,182	1,431	12,724	955	62,286
χ_q	109	480	123	2,329	267	4,146	447	16,905
χ_f	76	476	24	2,258	162	4,142	259	14,043
χ_i	8	7	15	235	103	2,429	105	4,107
$ \Sigma $	2.7	7.8	2.7	41.9	2.4	56.6	2.5	210.5

Table 1: Results for genuine and random signatures Σ

al. 2013) where two kinds of signatures are considered: *genuine signatures* corresponding to the signature of individual axioms, and *random signatures* with a given probability for a symbol to be included. For each type of signature and ontology, we took a sample of 400 runs and averaged module sizes. For random signatures we considered a probability of 1/1000. All experiments have been performed on a server with two Intel Xeon E5-2643 processors and 90GB of allocated RAM, running RDFox on 16 threads.

Table 1 summarises our results. We compared $\perp$ -modules with the modules for χ_c (Section 3.4) and $\top\perp^*$ -modules with those for χ_m , χ_q , χ_f , and χ_i (Section 3.3). We can see that module size consistently decreases as we consider weaker inseparability relations. In particular, the modules for χ_c can be 4 times smaller than $\perp$ -modules. The difference between $\top\perp^*$ -modules and χ_i modules is even bigger, especially in the case of FLY. In fact, χ_i modules are sometimes empty, which is not surprising since two predicates in a large ontology are unlikely to be in an implication relationship. Also note that our modules for semantic inseparability slightly improve on $\top\perp^*$ -modules. Finally, recall that our modules may not be minimal for their inseparability relation. Since techniques for extracting minimal modules are available only for model inseparability, and for restricted languages, we could not assess how close our modules are to minimal ones and hence the quality of our approximation.

Computation times were comparable for all settings χ^z with times being slightly higher for χ_i and χ_c as they involved a larger number of constants. Furthermore, extraction times were comparable to locality-based modules for genuine signatures with average times of 0.5s for FLY, 0.9s for SCT, 4.2 for GO and 5s for BM.

5 Conclusion and Future Work

We have proposed a novel approach to module extraction by exploiting off-the-shelf datalog reasoners, which allows us to efficiently compute approximations of minimal modules for different inseparability relations. Our results open the door to significant improvements in common applications of modules, such as computation of justifications, modular and incremental reasoning and ontology reuse, which currently rely mostly on locality-based modules.

Our approach is novel, and we see many interesting open problems. For example, the issue of optimality requires fur-

ther investigation. Furthermore, it would be interesting to integrate our extraction techniques in existing modular reasoners as well as in systems for justification extraction.

Acknowledgements

This work was supported by the Royal Society, the EPSRC projects MaSI³, Score! and DBOnto, and by the EU FP7 project OPTIQUE.

References

- Armas Romero, A.; Cuenca Grau, B.; and Horrocks, I. 2012. MORE: Modular combination of OWL reasoners for ontology classification. In *ISWC*, 1–16.
- Cuenca Grau, B.; Horrocks, I.; Kazakov, Y.; and Sattler, U. 2007. Just the right amount: Extracting modules from ontologies. In *WWW*, 717–726.
- Cuenca Grau, B.; Horrocks, I.; Kazakov, Y.; and Sattler, U. 2008. Modular reuse of ontologies: Theory and practice. *J. Artif. Intell. Res.* 31:273–318.
- Cuenca Grau, B.; Halaschek-Wiener, C.; Kazakov, Y.; and Suntisrivaraporn, B. 2010. Incremental classification of description logics ontologies. *J. Autom. Reason.* 44(4):337–369.
- Dantsin, E.; Eiter, T.; Gottlob, G.; and Voronkov, A. 2001. Complexity and expressive power of logic programming. *ACM Comput. Surv.* 33(3):374–425.
- Del Vescovo, C.; Parsia, B.; Sattler, U.; and Schneider, T. 2011. The modular structure of an ontology: Atomic decomposition. In *IJCAI*, 2232–2237.
- Del Vescovo, C.; Klinov, P.; Parsia, B.; Sattler, U.; Schneider, T.; and Tsarkov, D. 2013. Empirical study of logic-based modules: Cheap is cheerful. In *DL*, 144–155.
- Gatens, W.; Konev, B.; and Wolter, F. 2014. Lower and upper approximations for depleting modules of description logic ontologies. In *ECAI*, 345–350.
- Jiménez-Ruiz, E., and Cuenca Grau, B. 2011. LogMap: Logic-based and scalable ontology matching. In *ISWC*, 273–288.
- Jiménez-Ruiz, E.; Cuenca Grau, B.; Sattler, U.; Schneider, T.; and Berlanga Llavori, R. 2008. Safe and economic re-use of ontologies: A logic-based methodology and tool support. In *ESWC*, 185–199.
- Kalyanpur, A.; Parsia, B.; Horridge, M.; and Sirin, E. 2007. Finding all justifications of OWL DL entailments. In *ISWC*, 267–280.
- Konev, B.; Lutz, C.; Walther, D.; and Wolter, F. 2013. Model-theoretic inseparability and modularity of description logic ontologies. *Artif. Intell.* 203:66–103.
- Kontchakov, R.; Lutz, C.; Toman, D.; Wolter, F.; and Zakharyashev, M. 2011. The combined approach to ontology-based data access. In *IJCAI*, 2656–2661.
- Kontchakov, R.; Wolter, F.; and Zakharyashev, M. 2010. Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artif. Intell.* 174(15):1093–1141.
- Krötzsch, M.; Rudolph, S.; and Hitzler, P. 2008. ELP: Tractable rules for OWL 2. In *ISWC*, 649–664.
- Ludwig, M. 2014. Just: a tool for computing justifications w.r.t. $\mathcal{ELH}$ ontologies. In *ORE*, 1–7.
- Lutz, C., and Wolter, F. 2010. Deciding inseparability and conservative extensions in the description logic EL. *J. Symb. Comput.* 45(2):194–228.
- Marnette, B. 2009. Generalized schema-mappings: From termination to tractability. In *PODS*, 13–22.
- Motik, B.; Nenov, Y.; Piro, R.; Horrocks, I.; and Olteanu, D. 2014. Parallel materialisation of datalog programs in centralised, main-memory RDF systems. In *AAAI*, 129–137.
- Nortje, R.; Britz, K.; and Meyer, T. 2013. Reachability modules for the description logic *SRIQ*. In *LPAR*, 636–652.
- Sattler, U.; Schneider, T.; and Zakharyashev, M. 2009. Which kind of module should I extract? In *DL*.
- Seidenberg, J., and Rector, A. L. 2006. Web ontology segmentation: Analysis, classification and use. In *WWW*, 13–22.
- Stefanoni, G.; Motik, B.; and Horrocks, I. 2013. Introducing nominals to the combined query answering approaches for $\mathcal{EL}$. In *AAAI*, 1177–1183.
- Stuckenschmidt, H.; Parent, C.; and Spaccapietra, S., eds. 2009. *Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization*, volume 5445 of *LNCIS*. Springer.
- Suntisrivaraporn, B.; Qi, G.; Ji, Q.; and Haase, P. 2008. A modularization-based approach to finding all justifications for OWL DL entailments. In *ASWC*, 1–15.
- Suntisrivaraporn, B. 2008. Module extraction and incremental classification: A pragmatic approach for ontologies. In *ESWC*, 230–244.
- Tsarkov, D., and Palmisano, I. 2012. Chainsaw: A metareasoner for large ontologies. In *ORE*.
- Zhou, Y.; Nenov, Y.; Cuenca Grau, B.; and Horrocks, I. 2014. Pay-as-you-go OWL query answering using a triple store. In *AAAI*, 1142–1148.

A Inseparability Relations

We start by giving an alternative characterization of Σ -query and Σ -fact inseparability that will allow us to prove our results in a more uniform and clear way.

Proposition 22. *TBoxes $\mathcal{T}$ and $\mathcal{T}'$ are*

1. *Σ -query inseparable iff $\mathcal{T} \models r \Leftrightarrow \mathcal{T}' \models r$ holds for every rule r over Σ ;*
2. *Σ -fact inseparable iff $\mathcal{T} \models r \Leftrightarrow \mathcal{T}' \models r$ holds for every datalog rule r over Σ .*

Proof. It suffices to observe that, for every TBox $\mathcal{T}$ and every rule $r = \varphi \rightarrow \psi$ over Σ , $\mathcal{T} \models r$ iff $\mathcal{T} \cup \{\gamma\sigma \mid \gamma \in \varphi\} \models \psi\sigma$, with σ a substitution mapping all free variables in r to fresh, pairwise distinct constants. $\square$

Proposition 23. $\equiv^m \subsetneq \equiv^q \subsetneq \equiv^f \subsetneq \equiv^i$.

Proof. The inclusion $\equiv^q \subseteq \equiv^f$ is immediate by definition while $\equiv^f \subseteq \equiv^i$ follows by Proposition 22. The inclusion $\equiv^m \subseteq \equiv^q$ follows since $r^{\mathcal{I}} = r^{\mathcal{I}|_{\Sigma}} = r^{\mathcal{I}'|_{\Sigma}} = r^{\mathcal{I}'}$ for every rule r over Σ whenever $\mathcal{I}$ and $\mathcal{I}'$ coincide on Σ .

To show strictness of the inclusions, we can w.l.o.g. restrict ourselves to the signature $\Sigma = \{Q, \perp\}$ where Q is a unary predicate (if Σ contains more symbols, one can consider $\mathcal{T}$ such that $\text{Sig}(\mathcal{T}) \cap \Sigma \subseteq \{Q\}$; adapting the argument to higher arities for Q is also straightforward; finally, the presence of $\perp$ in Σ is not relevant for the proof).

For $\equiv^m \subsetneq \equiv^q$, suppose $\mathcal{T} = \{\top(x) \rightarrow \exists y.[R(x, y) \wedge A(y)], \top(x) \rightarrow \exists y.[R(x, y) \wedge B(y)], A(x) \wedge B(x) \rightarrow Q(x)\}$. Then $\mathcal{T} \equiv^q \emptyset$. However, $\mathcal{T} \not\equiv^m \emptyset$ since, for any interpretation $\mathcal{I}$ with a singleton domain such that $Q^{\mathcal{I}} = \emptyset$, $\mathcal{I}$ cannot be turned into a model of $\mathcal{T}$ without changing the interpretation of Q .

For $\equiv^q \subsetneq \equiv^f$, suppose $\mathcal{T} = \{\top(x) \rightarrow \exists y.[R(x, y) \wedge Q(y)]\}$. Then $\mathcal{T} \equiv^f \emptyset$ but $\mathcal{T} \not\equiv^q \emptyset$ since $\mathcal{T} \models \exists x.Q(x)$ while $\emptyset \not\models \exists x.Q(x)$.

For $\equiv^f \subsetneq \equiv^i$, suppose $\mathcal{T} = \{r\}$ where $r = Q(a) \wedge Q(b) \rightarrow Q(c)$. Then $\mathcal{T} \equiv^i \emptyset$ but $\mathcal{T} \not\equiv^f \emptyset$ since $\mathcal{T} \models r$ while $\emptyset \not\models r$. $\square$

B Deductive Inseparability

Theorems 3, 5, 7, 13 are all shown by a similar argument, which we present next.

Hyperresolution Given $r = \varphi(\mathbf{x}) \rightarrow \exists \mathbf{y}.[\bigvee_{i=1}^n \psi(\mathbf{x}, \mathbf{y})] \in \mathcal{T}$ we denote with $\text{sk}(r)$ the result of applying standard Skolemisation to r —which replaces, for each $y \in \mathbf{y}$, all occurrences of y in r by $f_y(\mathbf{x})$, where f_y is a fresh function symbol unique for y . Given a substitution θ mapping existentially quantified variables in $\mathcal{T}$ to constants and a Skolemised formula φ , we write $\Gamma_{\theta}(\varphi)$ for the formula obtained from φ by replacing every occurrence of a functional term $f_y(\mathbf{t})$ by the constant $y\theta$.

By distributing disjunctions over conjunctions in the head of $\text{sk}(r)$ we obtain a rule of the form $\varphi \rightarrow \bigwedge_{j=1}^m \psi'_j$ where each ψ'_j is a disjunction of atoms. We denote with $\text{cnf}(r)$ the set $\{\varphi \rightarrow \psi'_j \mid 1 \leq j \leq m\}$ and extend this notation in the natural way to $\text{cnf}(\mathcal{T}) = \bigcup_{r \in \mathcal{T}} \text{cnf}(r)$. We call $\text{cnf}(\mathcal{T})$ a *CNF TBox* and each $s \in \text{cnf}(\mathcal{T})$ a *CNF rule*. Clearly, $\text{cnf}(\mathcal{T}) \models \mathcal{T}$, and hence $\mathcal{T} \cup \mathcal{A} \models \varphi'$ implies $\text{cnf}(\mathcal{T}) \cup \mathcal{A} \models \varphi'$ for every $\mathcal{A}$ and φ' .

Let φ be a disjunction of facts, $\mathcal{A}$ an ABox, and $s = \bigwedge_{i=1}^n \gamma'_i \rightarrow \psi \in \text{cnf}(\mathcal{T})$. A formula φ is a *hyperresolvent* of s and ground disjunctions $\gamma_1 \vee \psi_1, \dots, \gamma_n \vee \psi_n$ (with each ψ_i potentially empty) if $\varphi = \bigvee_{i=1}^n \psi_i \vee \psi\sigma$ with σ a MGU of γ_i, γ'_i for each $1 \leq i \leq n$. Let $\mathcal{C}$ be a CNF TBox. A hyperresolution proof (or simply a *proof*) of φ in $\mathcal{C} \cup \mathcal{A}$ is a pair $\rho = (T, \lambda)$ where T is a tree, λ is a mapping from nodes in T to disjunctions of facts, and from edges in T to CNF rules in $\mathcal{C}$, such that for every node v the following properties hold:

1. $\lambda(v) = \varphi$ if v is the root of T ;
2. $\lambda(v) \in \mathcal{A}$ if v is a leaf in T ; and
3. if v has children $w_1, \dots, w_n$ then each edge from v to w_i is labelled by the same CNF rule s and $\lambda(v)$ is a hyperresolvent of s and $\lambda(w_1), \dots, \lambda(w_n)$.

If there exists a proof of φ in $\mathcal{C} \cup \mathcal{A}$ we write $\mathcal{C} \cup \mathcal{A} \vdash \varphi$. The *support* of φ is the set of CNF rules occurring in some proof of φ in $\mathcal{C} \cup \mathcal{A}$.

Hyperresolution is sound (if $\mathcal{C} \cup \mathcal{A} \vdash \varphi$ then $\mathcal{C} \cup \mathcal{A} \models \varphi$) and complete in the following sense: if $\mathcal{C} \cup \mathcal{A} \models \varphi$ then there exists $\psi \subseteq \varphi$ such that $\mathcal{C} \cup \mathcal{A} \vdash \psi$.

Given a module setting χ and $r \in \mathcal{T}$, we denote with $\Xi^{\chi}(r)$ the set of datalog rules in $\mathcal{P}^{\chi}$ corresponding to r , as described in Definition 1. The following auxiliary results provide the basis for correctness of our approach to module extraction.

Lemma 24. *Let $\chi = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ be a module setting. Let $\mathbf{N}$ be the set of constants mentioned in χ . Let $\mathcal{A}$ be a function-free ABox that only mentions constants that are fresh w.r.t. $\mathcal{T}$ and $\mathbf{N}$. Let ν be a mapping from constants in $\mathcal{A}$ to $\mathbf{N}$ such that $\mathcal{A}\nu \subseteq \mathcal{A}_0$. Let φ be a disjunction of facts and $\rho = (T, \lambda)$ a proof of φ in $\text{cnf}(\mathcal{T}) \cup \mathcal{A}$. The following properties hold:*

1. $\mathcal{P}^{\chi} \cup \mathcal{A}_0 \vdash \Gamma_{\theta}(\gamma\nu)$ for every $\gamma \in \varphi$.

2. For every $r \in \mathcal{T}$ such that ρ mentions some $s \in \text{cnf}(r)$ there exists $\gamma \in \varphi \cup \{\perp\}$ and a proof of $\Gamma_\theta(\gamma\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$ that mentions some rule in $\Xi^x(r)$.

Proof. We reason by induction on the depth d of ρ .

$d = 0$

In this case φ must be a fact in $\mathcal{A}$. Since $\mathcal{A}$ is function-free by assumption we have $\Gamma_\theta(\varphi\nu) = \varphi\nu$, and since $\mathcal{A}\nu \subseteq \mathcal{A}_0$ we have $\varphi\nu \in \mathcal{A}_0$. Therefore, there exists a trivial proof of $\Gamma_\theta(\varphi\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$ and property 1 is satisfied. Furthermore, if the depth of ρ is 0 then there cannot be any rules in its support, so property 2 is trivially satisfied as well.

$d > 0$

Let v be the root of T and $w_1, \dots, w_n$ the children of v . Then it must be

- $\lambda(w_i) = \delta_i \vee \psi_i$ for each $1 \leq i \leq n$;
- $\lambda(v, w_i) = s$ for each $1 \leq i \leq n$ with $s \in \text{cnf}(\mathcal{T})$ of the form $\bigwedge_{i=1}^n \delta'_i \rightarrow \varphi'$; and
- $\varphi = \bigvee_{i=1}^n \psi_i \vee \varphi'\sigma$ with σ a MGU of δ_i, δ'_i for each $1 \leq i \leq n$.

Consider $\gamma \in \varphi$. To show property 1 we need to find a proof of $\Gamma_\theta(\gamma\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$. If $\gamma \in \psi_i$ then by i.h. we can find such a proof. If $\gamma \in \varphi'\sigma$ then it must be $\gamma = \gamma'\sigma$ for some $\gamma' \in \varphi'$ and, by definition of $\mathcal{P}^x$, $\text{cnf}(\mathcal{T})$, and Γ_θ , $s \in \text{cnf}(\mathcal{T})$ implies $\bigwedge_{i=1}^n \delta'_i \rightarrow \Gamma_\theta(\gamma') \in \mathcal{P}^x$. Since σ is a MGU of δ_i, δ'_i (with $\delta_i = \delta'_i\sigma$) for each $1 \leq i \leq n$, $\sigma\nu$ must be a MGU of $\delta_i\nu, \delta'_i\nu$ (with $\delta_i\nu = \delta'_i\sigma\nu$) for each $1 \leq i \leq n$; furthermore, since δ'_i is necessarily function-free, it is $\Gamma_\theta(\delta'_i\sigma\nu) = \delta'_i\sigma\nu$, and thus $\Gamma_\theta(\delta_i\nu) = \delta'_i\sigma\nu$ and $\sigma\nu$ is also a MGU of $\Gamma_\theta(\delta_i\nu), \delta'_i$ for each $1 \leq i \leq n$. By i.h. we have a proof in $\mathcal{P}^x \cup \mathcal{A}_0$ of $\Gamma_\theta(\delta_i\nu)$ for each $1 \leq i \leq n$; it is easy to see that $\Gamma_\theta(\gamma')\sigma\nu = \Gamma_\theta(\gamma'\sigma\nu) = \Gamma_\theta(\gamma\nu)$, so combining these proofs with rule $\bigwedge_{i=1}^n \delta'_i \rightarrow \Gamma_\theta(\gamma')$ yields a proof of $\Gamma_\theta(\gamma\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$.

Now consider $r \in \mathcal{T}$ such that ρ mentions some $s' \in \text{cnf}(r)$. To show property 2 we need to find $\gamma \in \varphi \cup \{\perp\}$ and a proof of $\Gamma_\theta(\gamma\nu)$ that mentions some rule in $\Xi^x(r)$. Assume first that $s' = s$. If $\varphi' = \emptyset$ then it must be $\text{cnf}(r) = \{\bigwedge_{i=1}^n \delta'_i \rightarrow \perp\} \subseteq \Xi^x(r)$ and, as before, we can combine this rule with proofs in $\mathcal{P}^x \cup \mathcal{A}_0$ of the $\Gamma_\theta(\delta_i\nu)$ to obtain a proof of $\perp$ in $\mathcal{P}^x \cup \mathcal{A}_0$. If $\varphi' \neq \emptyset$ then it must be $\{\bigwedge_{i=1}^n \delta'_i \rightarrow \Gamma_\theta(\gamma') \mid \gamma' \in \varphi'\} \subseteq \Xi^x(r)$. Since $\varphi = \bigvee_{i=1}^n \psi_i \vee \varphi'\sigma$, for each $\gamma' \in \varphi'$ it is $\gamma'\sigma \in \varphi$ and, as we just saw, we can construct a proof of $\Gamma_\theta(\gamma'\sigma\nu)$ that mentions $\bigwedge_{i=1}^n \delta'_i \rightarrow \Gamma_\theta(\gamma')$. Finally, assume that $s' \neq s$. Then there must be some $i \in \{1, \dots, n\}$ such that s' is mentioned by the proof ρ_i of $\delta_i \vee \psi_i$ that is embedded in ρ . Since ρ_i is of depth $< d$, by i.h. there must be $\delta'' \in \delta_i \vee \psi_i$ and a proof ρ'' of $\Gamma_\theta(\delta''\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$ that mentions some rule in $\Xi^x(r)$. If $\delta'' \in \psi_i$ then $\delta'' \in \varphi$ already; if $\delta'' = \delta_i$ then, as before, for any $\gamma \in \varphi$ we can construct a proof of $\Gamma_\theta(\gamma\nu)$ in $\mathcal{P}^x \cup \mathcal{A}_0$ such that ρ'' is embedded in it. $\square$

Proposition 25. Let $r = \varphi(\mathbf{x}) \rightarrow \psi(\mathbf{x})$ with φ a conjunction and ψ a disjunction of atoms. Let $\mathcal{X} = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ be a module setting satisfying $\{\perp\} \subseteq \mathcal{A}_r$ and such that for every substitution σ mapping all variables in r to pairwise distinct constants not in $\mathcal{T}$ there exists a mapping ν_σ with $\varphi\sigma\nu_\sigma \subseteq \mathcal{A}_0$ and $\psi\sigma\nu_\sigma \subseteq \mathcal{A}_r$. Then

1. $\mathcal{T} \models r$ iff $\mathcal{M}^x \models r$;
2. if $\mathcal{T}' \subseteq \mathcal{T}$ is a justification for r in $\mathcal{T}$ then $\mathcal{T}' \subseteq \mathcal{M}^x$;
3. $\mathcal{T} \setminus \mathcal{M}^x \models r$ iff $\emptyset \models r$.

Proof.

1. By monotonicity, it is immediate that $\mathcal{T} \models r$ if $\mathcal{M}^x \models r$.

Suppose $\mathcal{T} \models r$ and let σ be a substitution mapping all variables in r to fresh, pairwise distinct constants. Then we have that $\mathcal{T} \cup \{\gamma\sigma \mid \gamma \in \varphi\} \models \psi\sigma$, which implies $\text{cnf}(\mathcal{T}) \cup \{\gamma\sigma \mid \gamma \in \varphi\} \models \psi\sigma$ and by completeness of hyperresolution $\text{cnf}(\mathcal{T}) \cup \{\gamma\sigma \mid \gamma \in \varphi\} \vdash \psi'$ for some $\psi' \subseteq \psi\sigma$. Since $\{\gamma\sigma\nu_\sigma \mid \gamma \in \varphi\} \subseteq \mathcal{A}_0$, by Lemma 24 we have that for each $s \in \mathcal{T}$ such that some $p \in \text{cnf}(s)$ supports ψ' in $\text{cnf}(\mathcal{T}) \cup \{\gamma\sigma \mid \gamma \in \varphi\}$ there exists $\gamma \in \Gamma_\theta(\psi'\nu_\sigma) \cup \{\perp\}$ that is supported in $\mathcal{P}^x \cup \mathcal{A}_0$ by some rule from $\Xi^x(s)$. By assumption, $\perp \in \mathcal{A}_r$; also, ψ' is function-free so $\Gamma_\theta(\psi'\nu_\sigma) = \psi'\nu_\sigma$, and hence, since $\psi\sigma\nu_\sigma \subseteq \mathcal{A}_r$, we have that $\psi'\nu_\sigma \subseteq \mathcal{A}_r$ and also $\gamma \in \mathcal{A}_r$. In either case we have $s \in \mathcal{M}^x$ and consequently $\mathcal{M}^x \models r$.

2. Let $\mathcal{T}' \subseteq \mathcal{T}$ be a justification for r in $\mathcal{T}$. As before, if σ is a ground substitution for r mapping variables in r to fresh, pairwise distinct constants, then $\text{cnf}(\mathcal{T}') \cup \{\gamma\sigma \mid \gamma \in \varphi\} \vdash \psi'$ for some $\psi' \subseteq \psi\sigma$. In fact, by minimality of justifications, for each $s \in \mathcal{T}'$ some $p \in \text{cnf}(s)$ must be in the support of some $\psi' \subseteq \psi\sigma$ in $\text{cnf}(\mathcal{T}') \cup \{\gamma\sigma \mid \gamma \in \varphi\}$. As before, by Lemma 24, this implies $s \in \mathcal{M}^x$.

3. By monotonicity, it is immediate that $\mathcal{T} \setminus \mathcal{M}^x \models r$ if $\emptyset \models r$.

Suppose $\mathcal{T} \setminus \mathcal{M}^x \models r$ and let $\mathcal{T}'$ be a justification for r in $\mathcal{T} \setminus \mathcal{M}^x$. Then $\mathcal{T}'$ is also a justification for r in $\mathcal{T}$, and, as we just proved, $\mathcal{T}' \subseteq \mathcal{M}^x$. This implies that $\mathcal{T}' = \emptyset$, and therefore $\emptyset \models r$. $\square$

Proposition 26. Let $r = \varphi(\mathbf{x}) \rightarrow \exists \mathbf{y}.\bigvee_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})$ be a rule. Let $\mathcal{X} = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ be a module setting satisfying

- $\{\perp\} \subseteq \mathcal{A}_r$ and also $\psi\sigma \subseteq \mathcal{A}_r$ for every substitution σ mapping all variables in r to constants in $\mathcal{X}$;

- for every substitution σ mapping all variables in r to pairwise distinct constants not in $\mathcal{T}$ there exists a mapping ν_σ such that $\varphi\sigma\nu_\sigma \subseteq \mathcal{A}_0$.

Then

1. $\mathcal{T} \models r$ iff $\mathcal{M}^x \models r$;
2. if $\mathcal{T}' \subseteq \mathcal{T}$ is a justification for r in $\mathcal{T}$ then $\mathcal{T}' \subseteq \mathcal{M}^x$;
3. $\mathcal{T} \setminus \mathcal{M}^x \models r$ iff $\emptyset \models r$.

Proof.

1. By monotonicity, it is immediate that $\mathcal{T} \cup \mathcal{A} \models r$ if $\mathcal{M}^x \cup \mathcal{A} \models r$.

Let Q be a fresh predicate and $\mathcal{T}_{\psi \rightarrow Q} = \{ \psi_i(\mathbf{x}, \mathbf{y}) \rightarrow Q(\mathbf{x}) \mid 1 \leq i \leq n \}$. Then

$$\mathcal{T} \models r \text{ iff } \mathcal{T} \cup \mathcal{T}_{\psi \rightarrow Q} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x}) \quad \text{and} \quad \mathcal{M}^x \models r \text{ iff } \mathcal{M}^x \cup \mathcal{T}_{\psi \rightarrow Q} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x})$$

Consider $\mathcal{T}' = \mathcal{T} \cup \mathcal{T}_{\psi \rightarrow Q}$ and $\Sigma' = \Sigma \cup \{Q\}$. Clearly, $\mathcal{T}'$ has the exact same existentially quantified variables as $\mathcal{T}$. Therefore $\mathcal{X}' = \langle \theta, \mathcal{A}_0, \mathcal{A}'_r \rangle$ with

$$\mathcal{A}'_r = \{ Q(\mathbf{x})\sigma \mid \sigma \text{ is a substitution mapping all variables in } \mathbf{x} \text{ to constants in } \mathcal{X} \}$$

is a module setting for $\mathcal{T}'$ and Σ' and by Proposition 25 we have that $\mathcal{T}' \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x})$ iff $\mathcal{M}^{\mathcal{X}'} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x})$.

If we show that $\mathcal{M}^{\mathcal{X}'} \setminus \mathcal{T}_{\psi \rightarrow Q} \subseteq \mathcal{M}^x$ then, by monotonicity, we will be able to conclude that

$$\mathcal{M}^{\mathcal{X}'} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x}) \text{ implies } \mathcal{M}^x \cup \mathcal{T}_{\psi \rightarrow Q} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x})$$

and thus that $\mathcal{T} \models r$ implies $\mathcal{M}^x \models r$.

Let $s \in \mathcal{M}^{\mathcal{X}'} \setminus \mathcal{T}_{\psi \rightarrow Q}$. Some $p \in \Xi^{\mathcal{X}'}(s) = \Xi^x(s)$ must be in the support of some $Q(\mathbf{x})\sigma \in \mathcal{A}'_r$ in $\mathcal{P}^{\mathcal{X}'} \cup \mathcal{A}_0$. In particular, p must be mentioned in some proof $\rho = (T, \lambda)$ of $Q(\mathbf{x})\sigma$ in $\mathcal{P}^{\mathcal{X}'} \cup \mathcal{A}_0$. Let v be the root of T and $w_1, \dots, w_m$ its children nodes, there must be some $\bigwedge_{j=1}^m \gamma_j(\mathbf{x}, \mathbf{y}) \rightarrow Q(\mathbf{x}) \in \mathcal{T}_{\psi \rightarrow Q}$ and a MGU σ' of $\gamma_j, \lambda(w_j)$ for each $1 \leq j \leq m$. Since $s \notin \mathcal{T}_{\psi \rightarrow Q}$, there must exist $j \in \{1, \dots, m\}$ such that p is mentioned in the proof ρ_j of $\lambda(w_j)$ in $\mathcal{P}^{\mathcal{X}'} \cup \mathcal{A}_0$ that is embedded in ρ . Furthermore, since Q does not occur in the body of any rule in $\mathcal{P}^{\mathcal{X}'} = \mathcal{P}^x \cup \mathcal{T}_{\psi \rightarrow Q}$, all rules mentioned in ρ_j must be in $\mathcal{P}^x$ and thus ρ_j is a proof of γ in $\mathcal{P}^x \cup \mathcal{A}_0$. Since by assumption $\lambda(w_j) = \gamma_j\sigma' \in \mathcal{A}_r$, this implies $s \in \mathcal{M}^x$.

2. Let $\mathcal{T}'' \subseteq \mathcal{T}$ be a justification for r in $\mathcal{T}$. As before, $\mathcal{T}'' \models r$ implies $\mathcal{T}'' \cup \mathcal{T}_{\psi \rightarrow Q} \models \varphi(\mathbf{x}) \rightarrow Q(\mathbf{x})$ and in particular for any substitution σ mapping variables in r to pairwise distinct constants we have $\mathcal{T}'' \cup \mathcal{T}_{\psi \rightarrow Q} \cup \{ \gamma\sigma \mid \gamma \in \varphi \} \models Q(\mathbf{x})\sigma$ and therefore $\text{cnf}(\mathcal{T}'' \cup \mathcal{T}_{\psi \rightarrow Q}) \cup \{ \gamma\sigma \mid \gamma \in \varphi \} \vdash Q(\mathbf{x})\sigma$. By minimality of justifications, for each $s \in \mathcal{T}''$ there must be some $p \in \text{cnf}(s)$ in the support of $Q(\mathbf{x})\sigma$ in $\text{cnf}(\mathcal{T}'' \cup \mathcal{T}_{\psi \rightarrow Q}) \cup \{ \gamma\sigma \mid \gamma \in \varphi \}$. It is easy to see that p must also be in the support of $Q(\mathbf{x})\sigma\nu_\sigma$ in $\text{cnf}(\mathcal{T}'' \cup \mathcal{T}_{\psi \rightarrow Q}) \cup \{ \gamma\sigma\nu_\sigma \mid \gamma \in \varphi \}$. Since $Q(\mathbf{x})\sigma\nu_\sigma \subseteq \mathcal{A}'_r$ and $\{ \gamma\sigma\nu_\sigma \mid \gamma \in \varphi \} \subseteq \mathcal{A}_0$, by Lemma 24 we have that $s \in \mathcal{M}^{\mathcal{X}'}$. In particular, since $s \in \mathcal{T}'' \subseteq \mathcal{T}$, it must be $s \in \mathcal{M}^{\mathcal{X}'} \setminus \mathcal{T}_{\psi \rightarrow Q} \subseteq \mathcal{M}^x$.
3. Again by monotonicity, it is immediate that $\mathcal{T} \setminus \mathcal{M}^x \cup \mathcal{A} \models r$ if $\mathcal{A} \models r$. By a similar argument to the one given in Proposition 25, it follows from 2 that any justification for r in $\mathcal{T} \setminus \mathcal{M}^x$ must be empty and therefore if $\mathcal{T} \setminus \mathcal{M}^x \models r$ then $\emptyset \models r$. $\square$

Theorem 3. $\mathcal{M}^{\mathcal{X}_i} \equiv_{\Sigma}^i \mathcal{T}$.

Proof. Consider an arbitrary rule of the form $A(\mathbf{x}) \rightarrow B(\mathbf{x})$ with $A, B \in \Sigma$ and $A \neq B$ (if $A = B$ the rule is tautological). Since $\mathbf{x}$ is implicitly universally quantified, we can assume w.l.o.g. that $\mathbf{x} = (x_1, \dots, x_n)$ with $x_1, \dots, x_n$ pairwise distinct. Let σ be a substitution mapping $x_1, \dots, x_n$, respectively, to $c_1, \dots, c_n$, pairwise distinct constants not in $\mathcal{T}$. Now consider a mapping ν_σ such that $c_i\nu_\sigma = c_A^i$. This mapping is well-defined because $c_1, \dots, c_n$ are pairwise distinct. By definition of $\mathcal{X}_i$, we have $A(\mathbf{x})\sigma\nu_\sigma \in \mathcal{A}_0^i$ and $B(\mathbf{x})\sigma\nu_\sigma \in \mathcal{A}_r^i$, and therefore, by Proposition 25, we have $\mathcal{T} \models A(\mathbf{x}) \rightarrow B(\mathbf{x})$ iff $\mathcal{M}^{\mathcal{X}_i} \models A(\mathbf{x}) \rightarrow B(\mathbf{x})$. $\square$

Theorem 5. $\mathcal{M}^{\mathcal{X}_f} \equiv_{\Sigma}^f \mathcal{T}$.

Proof. By Proposition 22 it suffices to show that for any datalog rule $r = \varphi \rightarrow \psi$ we have $\mathcal{T} \models r$ iff $\mathcal{M}^{\mathcal{X}_f} \models r$. Let σ be a substitution mapping all variables in r to pairwise distinct constants not in $\mathcal{T}$. Consider a mapping ν^* such that $x\sigma\nu^* = *$ for each $x \in \mathbf{x}$. Clearly $\varphi\sigma\nu^* \subseteq \mathcal{A}_0^f$ and $\psi\sigma\nu^* \subseteq \mathcal{A}_r^f$, and therefore, by Proposition 25, we have $\mathcal{T} \models r$ iff $\mathcal{M}^{\mathcal{X}_f} \models r$. $\square$

Theorem 7. $\mathcal{M}^{\mathcal{X}_q} \equiv_{\Sigma}^q \mathcal{T}$.

Proof. By Proposition 22 it suffices to show that for any rule $r = \varphi \rightarrow \psi$ we have $\mathcal{T} \models r$ iff $\mathcal{M}^{\mathcal{X}_q} \models r$. Let σ be a substitution mapping all variables in r to pairwise distinct constants not in $\mathcal{T}$. Given a mapping ν^* such that $x\sigma\nu^* = *$ for each $x \in \mathbf{x}$ it is clear that $\varphi\sigma\nu^* \subseteq \mathcal{A}_0^q$. It is also immediate that $\psi\sigma' \subseteq \mathcal{A}_r^q$ for every substitution σ' mapping all variables in r to constants in $\mathcal{X}_q$. Therefore, by Proposition 26, we have $\mathcal{T} \models r$ iff $\mathcal{M}^{\mathcal{X}_q} \models r$. $\square$

Theorem 13. $\mathcal{M}^{\mathcal{X}_c} \equiv_{\Sigma}^c \mathcal{T}$.

Proof. Analogous to the proof of Theorem 3. $\square$

C Model Inseparability

Given an ABox $\mathcal{A}$ and a datalog program $\mathcal{P}$, let $\mathcal{P}(\mathcal{A})$ denote the materialisation of $\mathcal{P} \cup \mathcal{A}$. Furthermore, given a module setting χ , let $\text{supp}(\chi)$ denote the support of χ .

Theorem 9. $\mathcal{M}^{\chi_m} \equiv_{\Sigma}^m \mathcal{T}$.

Proof. Let $\mathcal{I}$ be a model of $\mathcal{M}^{\chi_m}$. We assume w.l.o.g that $\mathcal{I}$ is defined over all of $\text{Sig}(\mathcal{T})$. Consider the interpretation $\mathcal{J}$ over $\text{Sig}(\mathcal{T})$ such that $\Delta^{\mathcal{J}} = \Delta^{\mathcal{I}}$ and

$$A^{\mathcal{J}} = \begin{cases} A^{\mathcal{I}} & \text{if } A \in (\Sigma \cup \text{Sig}(\text{supp}(\chi))) \setminus \{\perp\} \\ \Delta^{\text{arity}(A)} & \text{if } A \in \text{Sig}(\mathcal{P}^{\chi_m}(\mathcal{A}_0^m)) \setminus (\Sigma \cup \text{Sig}(\text{supp}(\chi))) \\ \emptyset & \text{otherwise} \end{cases}$$

Consider $r : \varphi(\mathbf{x}) \rightarrow \exists \mathbf{y}. [\bigvee_{j=1}^m \psi_j(\mathbf{x}, \mathbf{y})] \in \mathcal{T}$. We will show that $\mathcal{J} \models r$.

Assume first $m = 0$. Then $\Xi^{\chi_m}(r) = \{\varphi \rightarrow \perp\}$. If $r \in \mathcal{M}^{\chi_m}$ then in particular $\text{Sig}(r) \subseteq \text{Sig}(\text{supp}(\chi_m))$, so $\mathcal{I}$ and $\mathcal{J}$ agree over $\text{Sig}(r)$, and $\mathcal{J} \models r$. If $r \notin \mathcal{M}^{\chi_m}$ then, since $\perp \in \mathcal{A}_r^{\chi_m}$ and the only constant mentioned in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ is $*$, there must be $\gamma \in \varphi$ such that $\gamma* \notin \mathcal{P}^{\chi_m}(\mathcal{A}_0^m)$ (where, in an abuse of notation, $*$ denotes the substitution that maps all variables to $*$), and in particular $\text{Sig}(\gamma) \not\subseteq \text{Sig}(\mathcal{P}^{\chi_m}(\mathcal{A}_0^m))$. Since $\Sigma \cup \text{Sig}(\text{supp}(\chi_m)) \subseteq \text{Sig}(\mathcal{P}^{\chi_m}(\mathcal{A}_0^m))$, this implies that for $A \in \text{Sig}(\gamma)$ it is $A^{\mathcal{J}} = \emptyset$ and therefore trivially $\mathcal{J} \models r$.

Assume now $m > 0$ and let σ be a substitution over all variables in r such that $\mathcal{J} \models \varphi\sigma$ (if no such substitution exists then trivially $\mathcal{J} \models r$). Since $\Sigma \cup \text{Sig}(\text{supp}(\chi)) \subseteq \text{Sig}(\mathcal{P}^{\chi_m}(\mathcal{A}_0^m))$, all predicates in φ must occur in $\mathcal{P}^{\chi_m}(\mathcal{A}_0^m)$. In particular it must be $\gamma* \in \mathcal{P}^{\chi_m}(\mathcal{A}_0^m)$ for every $\gamma \in \varphi$. This implies $\delta* \in \mathcal{P}^{\chi_m}(\mathcal{A}_0^m)$ for every $\delta \in \bigcup_{j=1}^m \psi_j$ and therefore for every predicate A in $\text{Sig}(\bigvee_{j=1}^m \psi_j)$ we have that either $A^{\mathcal{J}} = A^{\mathcal{I}}$ or $A^{\mathcal{J}} = \Delta^{\text{arity}(A)}$ —in particular $A^{\mathcal{I}} \subseteq A^{\mathcal{J}}$. If $A^{\mathcal{J}} = \Delta^{\text{arity}(A)}$ for every $A \in \text{Sig}(\bigvee_{j=1}^m \psi_j)$, then it is immediate that $\mathcal{J} \models r$. Suppose there exists $A \in \text{Sig}(\bigvee_{j=1}^m \psi_j)$ such that $A^{\mathcal{J}} \neq \Delta^{\text{arity}(A)}$. Then $A \in \Sigma \cup \text{Sig}(\text{supp}(\chi_m))$. If $A \in \Sigma$ then $A(*, \dots, *) \in \mathcal{A}_r^{\chi_m}$. Since $A \in \text{Sig}(\bigvee_{j=1}^m \psi_j)$ and $\gamma* \in \mathcal{P}^{\chi_m}(\mathcal{A}_0^m)$ for every $\gamma \in \varphi$, there is a proof $\rho_{A,r}$ of $A(*, \dots, *)$ in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ that mentions a rule in $\Xi^{\chi_m}(r)$. Therefore $r \in \mathcal{M}^{\chi_m}$. If $A \in \text{Sig}(\text{supp}(\chi_m)) \setminus \Sigma$ then some other $\gamma' \in \mathcal{A}_r^{\chi_m}$ must be supported by a rule that has A in its signature. More specifically, there must be a proof of $\gamma' \in \mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ that has a proof of $A(*, \dots, *)$ as a subproof. Replacing this subproof with $\rho_{A,r}$ results in a proof of γ' in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ that mentions a rule in $\Xi^{\chi_m}(r)$. Therefore in this case $r \in \mathcal{M}^{\chi_m}$ too. Now, since all rules in $\Xi^{\chi_m}(r)$ have the same body as r , we have that $\text{Sig}(\varphi) \subseteq \text{supp}(\chi_m) \setminus \{\perp\}$ and therefore $\mathcal{I}$ and $\mathcal{J}$ agree over $\text{Sig}(\varphi)$. By assumption, $\mathcal{J} \models \varphi\sigma$, so also $\mathcal{I} \models \varphi\sigma$; furthermore $\mathcal{I} \models \mathcal{M}^{\chi_m}$ implies $\mathcal{I} \models \bigvee_{j=1}^m \psi_j\sigma$, which implies $\mathcal{J} \models \bigvee_{j=1}^m \psi_j\sigma$ because $A^{\mathcal{I}} \subseteq A^{\mathcal{J}}$ for every predicate $A \in \text{Sig}(\bigvee_{j=1}^m \psi_j)$. Since σ is arbitrary, we conclude that $\mathcal{J} \models r$. $\square$

D Depletingness and Preservation of Justifications

Theorem 15. $\mathcal{M}^{\chi_z}$ is depleting for each $z \in \{m, q, f, i, c\}$.

Proof. For $z \in \{q, f, i, c\}$ the statement follows from Propositions 22, 25 and 26 by the arguments already presented in the proofs for Theorems 3, 5 and 7.

For $z = m$, we will now show that $\mathcal{T} \setminus \mathcal{M}^{\chi_m} \equiv^m \emptyset$. Let $\mathcal{I}$ be a model of $\emptyset$ and $\mathcal{A}_{\mathcal{I}}$ the ABox defined by $\mathcal{I}$ over Σ . Consider the datalog program $\mathcal{P} = \bigcup_{r \in \mathcal{T} \setminus \mathcal{M}^{\chi_m}} \Xi_{\chi_m}(r)$, and the materialisation $\mathcal{P}(\mathcal{A}_{\mathcal{I}})$ of $\mathcal{P}$ w.r.t. $\mathcal{A}_{\mathcal{I}}$. We show that $\mathcal{P}(\mathcal{A}_{\mathcal{I}})$ is a model of $\mathcal{T} \setminus \mathcal{M}$ that coincides with $\mathcal{I}$ over Σ . For this, it suffices to show the following two properties:

- All facts over Σ in $\mathcal{P}(\mathcal{A}_{\mathcal{I}})$ must already be in $\mathcal{A}_{\mathcal{I}}$.
Let $\gamma \in \mathcal{P}(\mathcal{A}_{\mathcal{I}})$ be a fact over Σ . If $\gamma \notin \mathcal{A}_{\mathcal{I}}$ then there must exist a proof ρ of γ in $\mathcal{P} \cup \mathcal{A}_{\mathcal{I}}$. Since $\mathcal{A}_{\mathcal{I}}$ only mentions predicates from Σ and $\mathcal{P} \subseteq \mathcal{P}^{\chi_m}$, we can find a proof of $\gamma* \in \mathcal{A}_r^{\chi_m}$ in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ that mentions the exact same rules as ρ . Let r be a rule mentioned in ρ , there must exist $s \in \mathcal{T} \setminus \mathcal{M}^{\chi_m}$ such that $r \in \Xi^{\chi_m}(s)$; however, because r is also mentioned in a proof of $\gamma*$ in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$, it must also be $s \in \mathcal{M}^{\chi_m}$. This is a contradiction, so $\gamma \in \mathcal{A}_{\mathcal{I}}$.
- $\perp \notin \mathcal{P}(\mathcal{A}_{\mathcal{I}})$.
Suppose $\perp \in \mathcal{P}(\mathcal{A}_{\mathcal{I}})$. Then there must be a proof ρ of $\perp$ in $\mathcal{P} \cup \mathcal{A}_{\mathcal{I}}$. Again, we can find a proof of $\perp$ in $\mathcal{P}^{\chi_m} \cup \mathcal{A}_0^m$ supported by the exact same rules as ρ . Following a similar argument as before, we conclude that $\perp \notin \mathcal{P}(\mathcal{A}_{\mathcal{I}})$. $\square$

Theorem 17. Let $\mathcal{T}'$ be a justification for a first-order sentence φ in $\mathcal{T}$ and let $\text{Sig}(\varphi) \subseteq \Sigma$. Then, $\mathcal{T}' \subseteq \mathcal{M}^{\chi_m}$. Additionally, the following properties hold: (i) if φ is a rule, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_q}$; (ii) if φ is datalog, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_f}$; and (iii) if φ is of the form $A(\mathbf{x}) \rightarrow B(\mathbf{x})$, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_i}$; finally, if φ satisfies $A \in \Sigma, B \in \text{Sig}(\mathcal{T})$, then $\mathcal{T}' \subseteq \mathcal{M}^{\chi_c}$.

Proof. The claim follows from Propositions 22, 25 and 26 similarly to Theorems 3, 5 and 7. $\square$

E Module Containment

Definition 27. Let $\chi = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ and $\chi' = \langle \theta', \mathcal{A}'_0, \mathcal{A}'_r \rangle$ and let N and N' be the sets of constants mentioned in χ and χ' , respectively. A mapping $\mu : N \rightarrow N'$ is a *homomorphism from χ to χ'* if the following conditions hold: (i) $\theta' = \theta\mu$, (ii) $\mathcal{A}_0\mu \subseteq \mathcal{A}'_0$; and (iii) $\mathcal{A}_r\mu \subseteq \mathcal{A}'_r$. We write $\chi \hookrightarrow \chi'$ if a homomorphism from χ to χ' exists. $\diamond$

Theorem 28. If χ, χ' are s.t. $\chi \hookrightarrow \chi'$, then $\mathcal{M}^\chi \subseteq \mathcal{M}^{\chi'}$.

Proof. Suppose $\chi = \langle \theta, \mathcal{A}_0, \mathcal{A}_r \rangle$ and $\chi' = \langle \theta', \mathcal{A}'_0, \mathcal{A}'_r \rangle$ with N and N' the sets of constants mentioned in χ and χ' , respectively. Let μ be a homomorphism from χ to χ' and let $r \in \mathcal{M}^\chi$. Some $\gamma \in \mathcal{A}_r$ must be supported in $\mathcal{P}^\chi \cup \mathcal{A}_0$ by a rule in $\Xi^\chi(r)$. Since, by assumption, $\gamma\mu \in \mathcal{A}'_r$, it suffices for us to show that $\gamma\mu$ is supported in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ by a rule from $\Xi^{\chi'}(r)$.

To this end we will show that for any rule $r \in \mathcal{T}$ and any fact γ such that there exists a proof $\rho = (T, \lambda)$ of γ in $\mathcal{P}^\chi \cup \mathcal{A}_0$ mentioning $s \in \Xi^\chi(r)$, there exists a proof of $\gamma\mu$ in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ mentioning $\Xi^{\chi'}(r)$. We will reason by induction on the depth d of ρ —which must be at least 1 since by assumption it uses s .

$d = 1$

r must be of the form $\bigwedge_{i=1}^n \delta'_i(\mathbf{x}) \rightarrow \exists \mathbf{y} \cdot [\bigvee_{j=1}^m \psi_j(\mathbf{x}, \mathbf{y})]$ so s must be

– $\bigwedge_{i=1}^n \delta'_i \rightarrow \gamma'\theta$ with $\gamma' \in \psi_j$ for some $1 \leq j \leq m$ if $m > 0$

The λ -images of the leaves of T must be $\delta_1, \dots, \delta_n \in \mathcal{A}_0$ such that there exists a MGU σ of δ_i, δ'_i for every $1 \leq i \leq n$ satisfying $\gamma = \gamma'\theta\sigma$. By assumption, we have $\delta_i\mu \in \mathcal{A}'_0$ for every $1 \leq i \leq n$, and also $s' = \bigwedge_{i=1}^n \delta'_i \rightarrow \gamma'\theta' \in \Xi^{\chi'}(r)$ where $\theta' = \theta\mu$. Consider $\sigma' = \sigma\mu$. It is easy to see that $\mu\sigma\mu = \sigma\mu$ since the domain of σ is disjoint with both the domain and the range of μ . Therefore $(\delta_i\mu)\sigma\mu = \delta_i\sigma\mu$ for every $1 \leq i \leq n$. Furthermore, since σ is a MGU of δ_i, δ'_i for every $1 \leq i \leq n$, we have that σ' is a MGU of $\delta_i\mu, \delta'_i$ for every $1 \leq i \leq n$. Finally, since $\theta' = \theta\mu$, we have that $\gamma\mu = \gamma'\theta\sigma\mu = \gamma'\theta\mu\sigma\mu = \gamma'\theta'\sigma'$ is a consequence of s' and $\delta_1\mu, \dots, \delta_n\mu$, and hence we have a proof of $\gamma\mu$ in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ supported by $s' \in \Xi^{\chi'}(r)$.

– $\bigwedge_{i=1}^n \delta'_i \rightarrow \perp$ if $m = 0$

Then it must be $\gamma = \perp$ and, as in the previous case, the λ -images of the leaves of T must be $\delta_1, \dots, \delta_n \in \mathcal{A}_0$ such that there exists a MGU σ of δ_i, δ'_i for every $1 \leq i \leq n$. Also, $\delta_i\mu \in \mathcal{A}'_0$ for every $1 \leq i \leq n$, $s \in \Xi^\chi(r)$, and $\sigma' = \sigma\mu$ is a MGU of $\delta_i\mu, \delta'_i$ for every $1 \leq i \leq n$, so we have a proof of $\gamma\mu = \perp$ in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ supported by $s \in \Xi^{\chi'}(r)$.

$d > 1$

Let v be the root of T , let $\delta_1, \dots, \delta_n \in \mathcal{A}_0$ be the λ -images of the children of v and let $r' \in \mathcal{T}$ be such that the λ -image of the edges connecting v with its children in T is a rule in $\Xi^\chi(r')$. Either $s \in \Xi^\chi(r')$ or it is mentioned in some subproof of ρ . Our induction hypothesis implies that for each $r'' \in \mathcal{T}$, if some δ_i is supported in $\mathcal{P}^\chi \cup \mathcal{A}_0$ by a rule in $\Xi^\chi(r'')$, then also $\delta_i\mu$ is supported in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ by some rule in $\Xi^{\chi'}(r'')$. Therefore, in either case, following an argument similar to case $d = 1$, we can construct a proof ρ' of $\gamma\mu$ in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ from a collection of proofs of $\delta_1\mu, \dots, \delta_n\mu$ in $\mathcal{P}^{\chi'} \cup \mathcal{A}'_0$ and a rule in $\Xi^{\chi'}(r')$ in such a way that ρ' mentions a rule in $\Xi^{\chi'}(r)$. $\square$

Proposition 19. $\mathcal{M}^{\chi_i} \subseteq \mathcal{M}^{\chi_f} \subseteq \mathcal{M}^{\chi_a} \subseteq \mathcal{M}^{\chi_m} \subseteq \mathcal{M}^{\chi_b}$ and $\mathcal{M}^{\chi_i} \subseteq \mathcal{M}^{\chi_c} \subseteq \mathcal{M}^{\chi_b}$

Proof. This follows immediately from Theorem 28. $\square$

F Optimality

Definition 29. Let $\mathcal{T}$ be a TBox. Let $\Sigma \subseteq \text{Sig}(\mathcal{T})$ and $\Sigma' = \Sigma \setminus \{\perp\}$. For each existentially quantified variable y in $\mathcal{T}$, let c_y be a fresh constant. Let $\theta = \{y \mapsto c_y \mid y \text{ existentially quantified in } \mathcal{T}\}$. Furthermore, for each pair $\langle A, B \rangle \in \Sigma' \times \text{Sig}(\mathcal{T})$, let $\mathbf{c}_{A,B}$ be a vector of fresh constants of size $\text{arity}(A)$. We define $\Psi_0^i(\mathcal{T}, \Sigma) = \langle \theta, \mathcal{A}_0^i, \mathcal{A}_r^i \rangle$ where

- $\mathcal{A}_0^i = \{A(\mathbf{c}_{A,B}) \mid A, B \in \Sigma', A \neq B, \text{arity}(A) = \text{arity}(B)\} \cup \{A(\mathbf{c}_{A,\perp}) \mid A \in \Sigma'\}$
- $\mathcal{A}_r^i = \{B(\mathbf{c}_{A,B}) \mid A, B \in \Sigma', A \neq B, \text{arity}(A) = \text{arity}(B)\} \cup \{\perp\}$

We define $\Psi_0^c(\mathcal{T}, \Sigma) = \langle \theta, \mathcal{A}_0^c, \mathcal{A}_r^c \rangle$ where

- $\mathcal{A}_0^c = \{A(\mathbf{c}_{A,B}) \mid A \in \Sigma', B \in \text{Sig}(\mathcal{T}) \setminus \{\perp\}, A \neq B, \text{arity}(A) = \text{arity}(B)\} \cup \{A(\mathbf{c}_{A,\perp}) \mid A \in \Sigma'\}$
- $\mathcal{A}_r^c = \{B(\mathbf{c}_{A,B}) \mid A \in \Sigma', B \in \text{Sig}(\mathcal{T}) \setminus \{\perp\}, A \neq B, \text{arity}(A) = \text{arity}(B)\} \cup \{\perp\}$

For each predicate $B \in \Sigma$ and each $\mathbf{v} \in \{1, \dots, \text{arity}(B)\}^{\text{arity}(B)}$, let $*_{B,\mathbf{v}}^1, \dots, *_{B,\mathbf{v}}^{\text{arity}(B)+1}$ be fresh constants. We define $\Psi_0^f(\mathcal{T}, \Sigma) = \langle \theta, \mathcal{A}_0^f, \mathcal{A}_r^f \rangle$ where

- $\mathcal{A}_0^f = \{A(\mathbf{d}) \mid A \in \Sigma', B \in \Sigma, \mathbf{v} \in \{1, \dots, \text{arity}(B)\}^{\text{arity}(B)}, \mathbf{d} \in \{*_{B,\mathbf{v}}^1, \dots, *_{B,\mathbf{v}}^{\text{arity}(B)+1}\}^{\text{arity}(A)}, A(\mathbf{d}) \neq B(*_{B,\mathbf{v}}^{\mathbf{v}})\}$
- $\mathcal{A}_r^f = \{B(*_{B,\mathbf{v}}^{\mathbf{v}}) \mid B \in \Sigma, \mathbf{v} \in \{1, \dots, \text{arity}(B)\}^{\text{arity}(B)}\}$

$\diamond$

Proposition 30. *Let $z \in \{i, c\}$. Then, for every $\mathcal{T}$ and Σ , $\mathcal{M}^{\Psi^z(\mathcal{T}, \Sigma)} = \mathcal{M}^{\Psi_0^z(\mathcal{T}, \Sigma)}$.*

Proof. It is easy to see that $\Psi_0^z(\mathcal{T}, \Sigma) \hookrightarrow \Psi^z(\mathcal{T}, \Sigma)$ for $z \in \{i, c\}$. By Theorem 28, we thus have that $\mathcal{M}^{\Psi_0^z(\mathcal{T}, \Sigma)} \subseteq \mathcal{M}^{\Psi^z(\mathcal{T}, \Sigma)}$.

Before we continue, note that for each $z \in \{i, c\}$ the datalog programs $\mathcal{P}^{\Psi^z(\mathcal{T}, \Sigma)}$ and $\mathcal{P}^{\Psi_0^z(\mathcal{T}, \Sigma)}$ coincide. For readability, we will denote this program with $\mathcal{P}^z$.

Let $r \in \mathcal{M}^{\Psi^i(\mathcal{T}, \Sigma)}$. Some fact $\gamma \in \mathcal{A}_r^{x_i}$ must be supported by a rule in $\Xi^{\Psi^i(\mathcal{T}, \Sigma)}(r) = \Xi^{\Psi_0^i(\mathcal{T}, \Sigma)}(r)$. The fact γ must be either $\perp$ or $B(c_A)$ with $A, B \in \Sigma'$. It is easy to see how one can turn any proof of $\perp$ (resp. $B(c_A)$) in $\mathcal{P}^i \cup \mathcal{A}_0^i$ into a proof of $\perp$ (resp. $B(c_{A,B})$) in $\mathcal{P}^i \cup \mathcal{A}_0^i$ that mentions the exact same rules. By construction of Ψ^i , both $\perp$ and $B(c_{A,B})$ are in $\mathcal{A}_r^i$, so $r \in \mathcal{M}^{\Psi_0^i(\mathcal{T}, \Sigma)}$. Therefore $\mathcal{M}^{\Psi^i(\mathcal{T}, \Sigma)} \subseteq \mathcal{M}^{\Psi_0^i(\mathcal{T}, \Sigma)}$.

The argument for $z = c$ is analogous. $\square$

Theorem 21. Ψ^z is z -optimal for $z \in \{i, c\}$.

Proof. We show the claim for Ψ^i , the argument for Ψ^c is similar. Let $\Sigma' = \Sigma \setminus \{\perp\}$.

For Ψ^i , suppose for contradiction there is a uniform, i-admissible Ψ and some $\mathcal{T}$ such that $\mathcal{M}^{\Psi^i(\mathcal{T}, \Sigma)} \not\subseteq \mathcal{M}^{\Psi(\mathcal{T}, \Sigma)}$. Then, by Proposition 30, $\mathcal{M}^{\Psi_0^i(\mathcal{T}, \Sigma)} \not\subseteq \mathcal{M}^{\Psi(\mathcal{T}, \Sigma)}$, and hence, by Theorem 28, $\Psi_0^i(\mathcal{T}, \Sigma) \not\vdash \Psi(\mathcal{T}, \Sigma)$. Let $\Psi(\mathcal{T}, \Sigma) = \langle \theta', \mathcal{A}'_0, \mathcal{A}'_r \rangle$. Since $\Psi_0^i(\mathcal{T}, \Sigma) \not\vdash \Psi(\mathcal{T}, \Sigma)$, by construction of Ψ_0^i , there are two cases to consider:

- There are some $A, B \in \Sigma'$ with $\text{arity}(A) = \text{arity}(B)$ such that for every vector c of size $\text{arity}(A)$ of constants mentioned in Ψ , $A(c) \notin \mathcal{A}'_0$ or $B(c) \notin \mathcal{A}'_r$. Let

$$\mathcal{T}' = \{A(x) \rightarrow B(x)\} \cup \{\rightarrow \exists y. Q_y(y) \mid y \text{ existentially quantified in } \mathcal{T}, Q_y \text{ fresh for every } y\}$$

Then $\Psi(\mathcal{T}', \Sigma) = \Psi(\mathcal{T}, \Sigma)$ (by uniformity), and hence $\mathcal{M}^{\Psi(\mathcal{T}', \Sigma)} = \emptyset$. Since $\emptyset \not\models A(x) \rightarrow B(x)$, we have $\mathcal{M}^{\Psi(\mathcal{T}', \Sigma)} \not\models^i \mathcal{T}'$.

- We have $\perp \notin \mathcal{A}'_r$. Let

$$\mathcal{T}' = \{A(x) \rightarrow \perp\} \cup \{\rightarrow \exists y. Q_y(y) \mid y \text{ existentially quantified in } \mathcal{T}, Q_y \text{ fresh for every } y\}$$

for some $A \in \Sigma'$. Then $\Psi(\mathcal{T}', \Sigma) = \Psi(\mathcal{T}, \Sigma)$ (by uniformity), and hence $\mathcal{M}^{\Psi(\mathcal{T}', \Sigma)} = \emptyset$. Since $\emptyset \not\models A(x) \rightarrow \perp$, we have $\mathcal{M}^{\Psi(\mathcal{T}', \Sigma)} \not\models^i \mathcal{T}'$.

In both cases, we obtain a contradiction to Ψ being i-admissible. $\square$

Proposition 31. *The family Ψ_0^f is f-admissible.*

Proof. By Propositions 22 and 25, it suffices to show that, given a datalog rule $r = \varphi \rightarrow \gamma$ and a substitution σ mapping all variables in r to distinct constants, we can construct a mapping ν such that $\varphi\sigma\nu \subseteq \mathcal{A}_0^{f_0}$ and $\gamma\sigma\nu \in \mathcal{A}_r^{f_0}$. W.l.o.g. we can assume $\gamma \notin \varphi$ (otherwise r is a tautology and hence trivially entailed by $\mathcal{M}^{\Psi_0^f(\mathcal{T}, \Sigma)}$) and therefore $\gamma\sigma \notin \varphi\sigma$ by injectivity of σ .

Let $\gamma\sigma = B(c)$. We construct ν as follows. Let μ be an ordering of the constants in c . We define ν such that $c\nu = *_{B, c\mu}^{c\mu}$ if $c \in c$ and $c\nu = *_{B, c\mu}^{\text{arity}(B)+1}$ otherwise. Since $c\mu \in \{1, \dots, \text{arity}(B)\}^{\text{arity}(B)}$ we have $B(c)\nu \in \mathcal{A}_r^{f_0}$. Moreover, every fact in $\varphi\sigma$ is mapped by ν to a fact $A(d)$ where $A \in \Sigma \setminus \{\perp\}$, $d \in \{*_B^1, \dots, *_B^{\text{arity}(B)+1}\}^{\text{arity}(A)}$, and $A(d) \neq B(c)\nu$ since $B(c) \notin \varphi\sigma$. Thus $\varphi\sigma\nu \subseteq \mathcal{A}_r^{f_0}$, and the claim follows. $\square$

Proposition 32. *The family Ψ^f is not f-optimal.*

Proof. By Proposition 31, it suffices to show that $\mathcal{M}^{\Psi^f(\mathcal{T}, \Sigma)} \not\subseteq \mathcal{M}^{\Psi_0^f(\mathcal{T}, \Sigma)}$ for some $\mathcal{T}$ and Σ . Let $\mathcal{T} = \{A(x) \rightarrow A(x)\}$ and $\Sigma = \{A\}$. Then $\mathcal{M}^{\Psi^f(\mathcal{T}, \Sigma)} = \mathcal{T} \not\subseteq \emptyset = \mathcal{M}^{\Psi_0^f(\mathcal{T}, \Sigma)}$. $\square$