Efficient Axiomatization of OWL 2 EL Ontologies from Data by Means of Formal Concept Analysis

Francesco Kriegel

Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany
francesco.kriegel@tu-dresden.de

Abstract

We present an FCA-based axiomatization method that produces a complete OWL 2 EL TBox (the terminological part of an OWL 2 EL ontology) from a graph dataset in at most exponential time. We describe technical details that allow for efficient implementation as well as variations that dispense with the computation of extremely large axioms, thereby rendering the approach applicable albeit some completeness is lost. Moreover, we evaluate the prototype on real-world datasets.

Introduction

Description Logics (DLs) (Baader, Horrocks, Lutz, Sattler 2017) are formal languages used in knowledge-based systems that reason and make inferences about complex domains, particularly where precision and explainability are essential. By representing knowledge as ontologies built with DLs, these systems can perform automated reasoning to answer queries and thereby assist in making decisions based on the encoded knowledge. DLs are fundamental to the Semantic Web (Hitzler, Krötzsch, Rudolph 2010), a vision of the World Wide Web where information is represented in a machine-readable format. They provide the logical underpinning for the Web Ontology Language (OWL), which is widely used in the Semantic Web to enable better interoperability across different applications, domains, and natural languages.

In e-commerce, DL ontologies can be used to categorize products into different classes and sub-classes based on their attributes, features, and properties. This enables efficient search and navigation for users on e-commerce platforms, such as eBay and Alibaba (Shi, J. Chen, Dong, Khan, Liang, Zhou, Wu, Horrocks 2023). In finance, where accuracy and explicability are crucial, the DL formalism enables clear and unambiguous representation of financial concepts, such as assets, liabilities, investments, and transactions. Two examples are the Financial Regulation Ontology (FRO), and the Financial Industry Business Ontology (FIBO). Another example is the Dow Jones Knowledge Graph (Horrocks, Olivares, Cocchi, Motik, Roy 2022), which does not use DLs but similar Semantic Web technologies such as SHACL (Bogaerts, Jakubowski, Van den Bussche 2022).

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Moreover, DLs and ontologies have found extensive applications in healthcare and life sciences. The Systematized Nomenclature of Medicine – Clinical Terms (SNOMED CT, or SCT) is an ontology that represents medical terms used in electronic health records, such as clinical findings, symptoms, and diagnoses. It is employed in clinical decision support systems to assist healthcare professionals in making accurate diagnoses, suggesting appropriate treatments, and predicting outcomes based on patient-specific information. The Gene Ontology (GO), the world's largest source of information on the functions of genes, is a foundation for computational analysis of large-scale molecular biology and genetics experiments in biomedical research.

Among the different DLs, the $\mathcal{EL}$ family (Baader, Brandt, Lutz 2005, 2008) stands out as a lightweight option. $\mathcal{EL}$ is designed to strike a balance between expressivity and computational complexity, making it an ideal choice for applications where scalability and latency are crucial. It offers a more restricted set of constructs compared to other DLs, but can thus handle large-scale ontologies efficiently. The Web Ontology Language includes it as the profile OWL 2 EL.

Every DL ontology is subdivided into two parts. The ABox consists of factual statements about specific individuals or objects in the domain, such as assignment of individuals to concepts and linkage between individuals by roles. The TBox defines concepts and their hierarchy, roles and their characteristics, and constraints or rules that govern all individuals in the domain. By separating ABox and TBox, a DL ontology provides a clear distinction between instance-level and schema-level knowledge. This separation enables reusing the same TBox across different ABoxes, which in turn promotes scalability and maintainability, as changes in the TBox are propagated to all associated ABoxes.

SCT and GO are formulated in $\mathcal{EL}$. For instance, SCT contains the TBox statement

$$\begin{aligned} &\mathit{Common_cold} \\ &\quad \sqsubseteq \mathit{Disease} \sqcap \exists \mathit{causative_agent}.\mathit{Virus} \\ &\qquad \sqcap \exists \mathit{finding_site}.\mathit{Upper_respiratory_tract_structure} \\ &\qquad \sqcap \exists \mathit{pathological_process}.\mathit{Infectious_process} \end{aligned}$$

which is a concept inclusion (CI) and expresses that a common cold (the premise) is a disease that has as pathological process an infectious process, is caused by a virus, and can be found in the upper respiratory tract (the conclusion). We

could express that Alice is diagnosed with common cold by the ABox statement $alice : \exists has_diagnose.Common_cold$.

Building and maintaining DL ontologies is a laborious task, especially for large domains. Knowledge engineers and domain experts work together to transfer their knowledge into an ontology. While the ABox is usually filled with observed data, constructing the TBox is a more complex endeavour. Assistance by automated approaches or guidance by interactive approaches is often valuable. For instance, a selection of individuals in the data can be described by a single concept (Funk, Jung, Lutz 2022; Funk, Jung, Lutz, Pulcini, Wolter 2019; Zarrieß, Turhan 2013) that the experts integrate into the conceptual hierarchy of the TBox. They can also model the schema of an ontology as a diagram (similar to UML class diagrams) that is then automatically translated into a TBox (Sarker, Krisnadhi, Hitzler 2016). New ontologies can be constructed from existing ones as well. Two or more ontologies can be integrated by ontology alignment (J. Chen, Jiménez-Ruiz, Horrocks, X. Chen, Myklebust 2023; Jimeno-Yepes, Jiménez-Ruiz, Llavori, Rebholz-Schuhmann 2009). Conversely, a part of an ontology representing a sub-domain can be extracted by modularization (Cuenca Grau, Horrocks, Kazakov, Sattler 2008), uniform interpolation or forgetting (Lutz, Wolter 2011; Zhao, Schmidt, Wang, Zhang, Feng 2020), or other techniques (Alghamdi, Schmidt, Del-Pinto, Gao 2021).

Formal Concept Analysis (FCA) (Ganter, Wille 1999) is a mathematical theory that represents data as formal contexts in which objects are described by their attributes. These attributes are similar to atomic statements in propositional logic and unary predicates in first-order logic. FCA has two main applications: the concept lattice reveals the conceptual hierarchy in the data (Wille 1982), and the canonical implication base is a complete set of implications, i.e. it entails all implications valid in the data (Guigues, Duquenne 1986; Stumme 1996). No complete set with fewer implications exists (Distel 2011; Wild 1994). If the data is not explicitly available but is only known by an expert, attribute exploration (Ganter 1984) enables interactive construction of the implication base.

The data analysis capabilities of FCA have been successfully employed in DLs, especially for the construction and extension of DL ontologies. Given a finite set of concepts, the hierarchy of all their conjunctions can efficiently be computed (Baader 1995). In order to support a bottom-up construction of DL ontologies, one can first compute most specific concepts for all individuals and then efficiently build the hierarchy of their least common subsumers (Baader, Molitor 2000; Baader, Sertkaya 2004). An ontology should be extended when it is incomplete since missing statements have been identified that should be entailed. Interactively completing the ontology using FCA is possible when attention is restricted to CIs over conjunctions from a fixed set of concepts (Baader, Ganter, Sertkaya, Sattler 2007). Moreover, it can be extended with new statements that are guessed by machine-learning approaches based on knowledge-graph embeddings (Jackermeier, J. Chen, Horrocks 2023; Shi, J. Chen, Dong, Khan, Liang, Zhou, Wu, Horrocks 2023). However, some of these embeddings fail to capture the semantics (Jain, Kalo, Balke, Krestel 2021) and, in effect, a large amount of useless, false predictions might be generated. This major issue can possibly be remedied by novel embedding approaches (Abboud, Ceylan, Lukasiewicz, Salvatori 2020; Asaadi, Giesbrecht, Rudolph 2023).

Axiomatization is another approach to constructing ontologies. In general, axiomatization is the task of describing a dataset (or any other formal object) by means of logical statements or axioms, viz. such that a logical formula (in the underlying logic) holds in the data iff. that formula is entailed by these axioms. In addition, axiomatization enables data analysis by transferring the given data into meaningful logical statements. By a suitable choice of the logical formalism, interesting and condensed insights into the analyzed data can be obtained.

In FCA, the canonical implication base axiomatizes data in form of a formal context by means of implications in propositional logic. By exploiting the similarity between $\mathcal{EL}$ CIs and FCA implications, a complete $\mathcal{EL}$ TBox can be axiomatized from observed graph data (Baader, Distel 2008). If the data is deemed incomplete, the latter approach can interact with the experts to ask for additional data when the validity of a TBox statement cannot be determined yet (Baader, Distel 2009). Both the unsupervised and the interactive approach terminate with a TBox that is sound and complete for the provided data, i.e. it entails a TBox statement if and only if that statement holds in the data. Moreover, the FCA-based axiomatization method was extended towards more expressive DLs (Kriegel 2017, 2019). Confident CI bases axiomatize all CIs that are valid already for a sufficiently large portion of all objects (Borchmann 2013, 2015). There are also other interactive approaches (Klarman, Britz 2015; Konev, Lutz, Ozaki, Wolter 2017) but which seem to have only limited practical value since the experts are required to terminate the process manually when they believe that the target TBox has been found (i.e. completeness of the constructed TBox is not guaranteed by the approach, but must be detected by the experts).

Our contributions are as follows. We reconsider the FCA-based approach to completely axiomatizing $\mathcal{EL}$ CIs from graph data (Baader, Distel 2008, 2009) and

1. thoroughly revise and simplify its technical description including proofs,
2. equip it with support for already known CIs valid in the data (thus enabling it for ontology completion),
3. analyze its computational complexity,
4. explain how further types of TBox statements supported by the $\mathcal{EL}$ family that are not just syntactic sugar can be completely axiomatized, viz. range restrictions and role inclusions,
5. describe how it can be implemented efficiently,
6. introduce variations that dispense with the computation of disjointness axioms or extremely large CIs without practical relevance, thereby rendering the approach applicable in practise, albeit some completeness is lost,
7. and evaluate the implementation on real-world datasets.

For space restrictions, technical details and proofs can be found in the extended version (Kriegel 2023a).

Preliminaries

The $\mathcal{EL}$ Family of DLs and OWL 2 EL

Fix a signature consisting of *individual names (INs)*, *concept names (CNs)*, and *role names (RNs)*. *Concept descriptions (CDs)* are built by $C ::= \top \mid \bot \mid A \mid C \sqcap C \mid \exists r.C$ where A ranges over all CNs and r over all RNs. A *TBox* is a finite set of *concept inclusions (CIs)* $C \sqsubseteq D$, *range restrictions (RRs)* $\top \sqsubseteq \forall r.C$, and *role inclusions (RIs)* $R \sqsubseteq s$, involving CDs C, D, RNs r, s, and *role chains* $R ::= \varepsilon \mid r \mid R \circ R$. An *ABox* is a finite set of *concept assertions (CAs)* $a{:}C$ and *role assertions (RAs)* $(a, b) : r$. An *ontology* consists of a TBox and an ABox. The $\mathcal{EL}$ family and OWL 2 EL additionally allow for nominals $\{a\}$ in CDs, but we ignore these to avoid overfitting in the axiomatization method. We also ignore concrete domains (datatypes for strings, numbers, etc.) as no $\mathcal{EL}$ reasoner currently supports them. As syntactic sugar we have *disjointness axioms* $C_1 \sqcap \cdots \sqcap C_n \sqsubseteq \bot$, *domain restrictions* $\exists r.\top \sqsubseteq C$, *concept equivalences* $C \equiv D$, *role equivalences* $r \equiv s$, *transitivity axioms* $r \circ r \sqsubseteq r$, and *reflexivity axioms* $\varepsilon \sqsubseteq r$.

$\mathcal{EL}$ can be translated into first-order logic and thus has a model-theoretic semantics, based on *interpretations* $\mathcal{I}$ consisting of a non-empty set $\mathsf{Dom}(\mathcal{I})$, called the *domain*, and of a function $\cdot^{\mathcal{I}}$ that gives meaning to the INs, CNs, and RNs.

Reasoning is the process of deciding or enumerating consequences of an ontology. An ontology $\mathcal{O}$ *entails* an axiom α, written $\mathcal{O} \models \alpha$, if α is satisfied in every model of $\mathcal{O}$. In this case α follows from the axioms in $\mathcal{O}$ by logical inference. Entailment in the $\mathcal{EL}$ family can be decided in polynomial time with the Completion algorithm (Baader, Brandt, Lutz 2005, 2008), which uses rules to materialize consequences and is implemented in the reasoner ELK (Kazakov, Klinov 2015; Kazakov, Krötzsch, Simančik 2014). Subsumption is a special form of entailment: D *subsumes* C w.r.t. $\mathcal{T}$, written $C \sqsubseteq^{\mathcal{T}} D$, if $\mathcal{T}$ entails $C \sqsubseteq D$.

Simulations and the DL $\mathcal{EL}^{\bot}_{\mathsf{si}}$

Given interpretations $\mathcal{I}$ and $\mathcal{J}$, a *simulation* from $\mathcal{I}$ to $\mathcal{J}$ is a relation $\mathfrak{S} \subseteq \mathsf{Dom}(\mathcal{I}) \times \mathsf{Dom}(\mathcal{J})$ such that **(S1)** if $x \in A^{\mathcal{I}}$ and $(x, y) \in \mathfrak{S}$, then $y \in A^{\mathcal{J}}$, and **(S2)** if $(x, x') \in r^{\mathcal{I}}$ and $(x, y) \in \mathfrak{S}$, then $(y, y') \in r^{\mathcal{J}}$ and $(x', y') \in \mathfrak{S}$ for some y'. We write $(\mathcal{I}, x) \preceq (\mathcal{J}, y)$ or just $x \preceq y$ if there is a simulation from $\mathcal{I}$ to $\mathcal{J}$ that contains (x, y).

Since the empty relation is a simulation and the union of simulations is a simulation, there is a maximal simulation from $\mathcal{I}$ to $\mathcal{J}$, denoted by $\mathfrak{S}_{\mathcal{I},\mathcal{J}}$. It is computable in polynomial time by starting with the full relation and subsequently deleting pairs that violate Condition (S1) or (S2). The relations $\preceq$ and $\mathfrak{S}_{\mathcal{I},\mathcal{J}}$ are equal since $x \preceq y$ iff. $(x, y) \in \mathfrak{S}_{\mathcal{I},\mathcal{J}}$.

Simulations characterize the semantics of $\mathcal{EL}$: we can rewrite each CD C but $\bot$ into a tree-shaped interpretation $\mathcal{I}_C$ with root x_C such that $C^{\mathcal{J}} = \{\, y \mid (\mathcal{I}_C, x_C) \preceq (\mathcal{J}, y) \,\}$ for each $\mathcal{J}$. The idea underlying the DL $\mathcal{EL}^{\bot}_{\mathsf{si}}$ (Lutz, Piro, Wolter 2010) is to replace $(\mathcal{I}_C, x_C)$ with an arbitrary finite, pointed interpretation $(\mathcal{I}, x)$: it features the additional CDs $\exists^{\mathsf{sim}}(\mathcal{I}, x)$ where $(\exists^{\mathsf{sim}}(\mathcal{I}, x))^{\mathcal{J}} := \{\, y \mid (\mathcal{I}, x) \preceq (\mathcal{J}, y) \,\}$. $\mathcal{EL}_{\mathsf{si}}$ and $\mathcal{EL}_{\mathsf{gfp}}$ (Baader 2003) are polynomially equivalent.

We often use denotations $\exists^{\mathsf{sim}}(\mathcal{C}, c)$, $\exists^{\mathsf{sim}}(\mathcal{D}, d)$, ... for such CDs. For convenience we may use ABox notation: we specify $\mathsf{Dom}(\mathcal{C})$ as usual but describe the function $\cdot^{\mathcal{C}}$ by the set $\{\, x : A \mid x \in A^{\mathcal{C}} \,\} \cup \{\, (x, y) : r \mid (x, y) \in r^{\mathcal{C}} \,\}$.

Extensions of $\mathcal{EL}^{\bot}_{\mathsf{si}}$ CDs can be read off from maximal simulations as follows, where $\mathfrak{S}(x) := \{\, y \mid (x, y) \in \mathfrak{S} \,\}$.

Lemma 1. $(\exists^{\mathsf{sim}}(\mathcal{C}, c))^{\mathcal{I}} = \mathfrak{S}_{\mathcal{C},\mathcal{I}}(c)$

Lemma 2. $\exists^{\mathsf{sim}}(\mathcal{C}, c) \sqsubseteq^{\emptyset} \exists^{\mathsf{sim}}(\mathcal{D}, d)$ *iff.* $(\mathcal{D}, d) \preceq (\mathcal{C}, c)$

Given $C := \exists^{\mathsf{sim}}(\mathcal{C}, c)$, the set of its *top-level conjuncts* is $\mathsf{Conj}(C) := \{\, A \mid c \in A^{\mathcal{C}} \,\} \cup \{\, \exists r.\exists^{\mathsf{sim}}(\mathcal{C}, u) \mid (c, u) \in r^{\mathcal{C}} \,\}$, and $C{\restriction}_n$ denotes its *unfolding* into an $\mathcal{EL}$ CD up to depth n.

We introduce a rule-based calculus with which subsumption w.r.t. a TBox $\mathcal{T}$ of $\mathcal{EL}^{\bot}_{\mathsf{si}}$ CIs can be decided.

$\sqsubseteq_{+}$-Rule. If there is an object x in $\mathsf{Dom}(\mathcal{C})$ and a CI $\exists^{\mathsf{sim}}(\mathcal{E}, e) \sqsubseteq \exists^{\mathsf{sim}}(\mathcal{F}, f)$ in $\mathcal{T}$ such that $(\mathcal{E}, e) \preceq (\mathcal{C}, x)$ but $(\mathcal{F}, f) \not\preceq (\mathcal{C}, x)$, then yield the interpretation $\mathcal{C}'$ with domain $\mathsf{Dom}(\mathcal{C}') := \mathsf{Dom}(\mathcal{C}) \cup \mathsf{Dom}(\mathcal{F})$ and function $\cdot^{\mathcal{C}'} := \cdot^{\mathcal{C}} \cup \cdot^{\mathcal{F}} \cup \{\, x : A \mid f \in A^{\mathcal{F}} \,\} \cup \{\, (x, y) : r \mid (f, y) \in r^{\mathcal{F}} \,\}$.

$\sqsubseteq_{\bot}$-Rule. If there is an object x in $\mathsf{Dom}(\mathcal{C})$ and a CI $\exists^{\mathsf{sim}}(\mathcal{E}, e) \sqsubseteq \bot$ in $\mathcal{T}$ such that $(\mathcal{E}, e) \preceq (\mathcal{C}, x)$, then fail.

Consider a CD $C := \exists^{\mathsf{sim}}(\mathcal{C}, c)$. Exhaustively applying both rules either fails in which case C is unsatisfiable w.r.t. $\mathcal{T}$ and we define $C^{\mathcal{T}} := \bot$, or produces an interpretation $\mathcal{C}'$ and then we set $C^{\mathcal{T}} := \exists^{\mathsf{sim}}(\mathcal{C}', c)$. Moreover, let $\bot^{\mathcal{T}} := \bot$. We call $C^{\mathcal{T}}$ the *most specific consequence* of C w.r.t. $\mathcal{T}$.

Proposition 3. *Subsumption in $\mathcal{EL}^{\bot}_{\mathsf{si}}$ can be decided in polynomial time. In particular, $C \sqsubseteq^{\mathcal{T}} D$ iff. $C^{\mathcal{T}} \sqsubseteq^{\emptyset} D$.*

Formal Concept Analysis

FCA is concerned with analyzing a *formal context* $\mathbb{K} := (G, M, I)$, where G is a set of *objects*, M is a set of *attributes*, and $I \subseteq G \times M$ is an *incidence relation*. We express by $(g, m) \in I$ that the object g *has* the attribute m. An *implication* is an expression $U \to V$ where U and V are subsets of M. The context $\mathbb{K}$ *satisfies* $U \to V$ if every object in G that has all attributes in U also has each attribute in V; and $\mathbb{K}$ is a *model* of an implication set $\mathcal{L}$ if it satisfies every implication in $\mathcal{L}$. Furthermore, $\mathcal{L}$ *entails* $U \to V$, denoted as $\mathcal{L} \models U \to V$, if $U \to V$ is satisfied in every model of $\mathcal{L}$. Implication entailment can be decided in linear time by means of the algorithm LinClosure (Beeri, Bernstein 1979).

An *implication base* of $\mathbb{K}$ relative to $\mathcal{L}$ is an implication set $\mathcal{B}$ of which $\mathbb{K}$ is a model and that together with $\mathcal{L}$ is *complete*, i.e. $\mathcal{B} \cup \mathcal{L}$ entails each implication satisfied in $\mathbb{K}$. The *canonical implication base* $\mathsf{Can}(\mathbb{K}, \mathcal{L})$ can be computed in exponential time with the algorithm NextClosure (Ganter 1984), and no base with fewer implications exists.

FCA can be seen as $\mathcal{EL}$ without RNs and $\bot$. In particular, each formal context is an interpretation without RNs, and implications $U \to V$ are CIs $\bigsqcap U \sqsubseteq \bigsqcap V$, using the syntactic sugar $\bigsqcap\{C_1, \ldots, C_n\} := C_1 \sqcap \cdots \sqcap C_n$ and $\bigsqcap \emptyset := \top$.

Axiomatization of $\mathcal{EL}$ TBoxes

We first focus on axiomatizing CIs. RRs and RIs will be considered later. As input we expect graph data in form of an interpretation $\mathcal{I}$, which includes knowledge graphs, graph databases, and RDF data: the CNs are the node labels and the RNs are the edge labels. Preprocessing of a

knowledge graph might be necessary, e.g. to correctly treat the metadata as well as to materialize the modelling conventions (Krötzsch 2019). Further given is a TBox $\mathcal{T}$ of which $\mathcal{I}$ is a model and that contains known CIs that should be preserved by the axiomatization. Note that $\mathcal{T}$ might be empty. We will compute a CI base in the following sense.

Definition 4. A TBox is *complete* for $\mathcal{I}$ if it entails all CIs satisfied in $\mathcal{I}$. A *CI base* of $\mathcal{I}$ relative to $\mathcal{T}$ is a TBox $\mathcal{B}$ of which $\mathcal{I}$ is a model and such that $\mathcal{B} \cup \mathcal{T}$ is complete for $\mathcal{I}$.

A CI base $\mathcal{B}$ together with the given TBox $\mathcal{T}$ axiomatizes all CIs satisfied in $\mathcal{I}$. We also call $\mathcal{B}$ a *completion* of $\mathcal{T}$ w.r.t. $\mathcal{I}$ as we obtain a complete TBox by adding all CIs in $\mathcal{B}$ to $\mathcal{T}$.

We will convert the interpretation $\mathcal{I}$ into a formal context such that its implication base can be rewritten into a CI base of $\mathcal{I}$. We therefore use induced contexts (Rudolph 2004).

Definition 5. Let $\mathbf{M}$ be a set of CDs. The *induced context* is $\mathbb{K}_\mathcal{I} := (\mathsf{Dom}(\mathcal{I}), \mathbf{M}, I)$ where $(x, C) \in I$ iff. $x \in C^\mathcal{I}$.

Lemma 6. *Given subsets $\mathbf{C}, \mathbf{D} \subseteq \mathbf{M}$, the CI $\bigsqcap \mathbf{C} \sqsubseteq \bigsqcap \mathbf{D}$ is satisfied in $\mathcal{I}$ iff. the implication $\mathbf{C} \to \mathbf{D}$ is satisfied in $\mathbb{K}_\mathcal{I}$.*

It follows that, if $\mathcal{B}$ is an implication base of $\mathbb{K}_\mathcal{I}$, then the TBox $\bigsqcap \mathcal{B} := \{ \bigsqcap \mathbf{C} \sqsubseteq \bigsqcap \mathbf{D} \mid \mathbf{C} \to \mathbf{D} \in \mathcal{B} \}$ has $\mathcal{I}$ as a model. Whether this TBox is complete depends on the choice of the attribute set $\mathbf{M}$, which we will address next.

Model-based Most Specific Concepts

For each subset X of $\mathsf{Dom}(\mathcal{I})$, we denote by $X^\mathcal{I}$ the *model-based most specific CD (MMSCD)* of X in $\mathcal{I}$ that is determined up to equivalence by **(M1)** $X \subseteq (X^\mathcal{I})^\mathcal{I}$, and **(M2)** for each CD C, if $X \subseteq C^\mathcal{I}$, then $X^\mathcal{I} \sqsubseteq^\emptyset C$. We will omit braces and write $X^{\mathcal{I}\mathcal{I}}$ instead of $(X^\mathcal{I})^\mathcal{I}$.

As in FCA the extended interpretation function $C \mapsto C^\mathcal{I}$ and the MMSCD mapping $X \mapsto X^\mathcal{I}$ form a Galois connection, i.e. all subsets X, Y of $\mathsf{Dom}(\mathcal{I})$ and all CDs C, D satisfy the following properties (Baader, Distel 2008).

(G1) $X \subseteq C^\mathcal{I}$ iff. $X^\mathcal{I} \sqsubseteq^\emptyset C$

(G2) $X^\mathcal{I} \sqsubseteq^\emptyset Y^\mathcal{I}$ if $X \subseteq Y$ **(G5)** $C^\mathcal{I} \subseteq D^\mathcal{I}$ if $C \sqsubseteq^\emptyset D$

(G3) $X \subseteq X^{\mathcal{I}\mathcal{I}}$ **(G6)** $C^{\mathcal{I}\mathcal{I}} \sqsubseteq^\emptyset C$

(G4) $X^\mathcal{I} \equiv^\emptyset X^{\mathcal{I}\mathcal{I}\mathcal{I}}$ **(G7)** $C^\mathcal{I} = C^{\mathcal{I}\mathcal{I}\mathcal{I}}$

For cycles in the interpretation $\mathcal{I}$, some MMSCDs might not be expressible in $\mathcal{EL}^\bot$ but in $\mathcal{EL}^\bot_{\mathsf{si}}$. The MMSCD of $\emptyset$ is $\bot$ and the MMSCD of each singleton $\{x\}$ is $\exists^{\mathsf{sim}}(\mathcal{I}, x)$. MMSCDs of sets with two or more objects can be efficiently computed with the powering, which is a permutation-invariant representation of all powers of $\mathcal{I}$ of any arity.

Definition 7. The *powering* $\wp(\mathcal{I})$ is the interpretation with domain $\mathsf{Dom}(\wp(\mathcal{I})) := \wp(\mathsf{Dom}(\mathcal{I}))$ and interpr. function $\cdot^{\wp(\mathcal{I})}$ where $X \in A^{\wp(\mathcal{I})}$ if $X \subseteq A^\mathcal{I}$, and $(X, Y) \in r^{\wp(\mathcal{I})}$ if Y is a minimal hitting set of $\{ r^\mathcal{I}(x) \mid x \in X \}$.

Recall that a *hitting set* of a set $\mathcal{S}$ of sets is a set H such that $H \cap S \neq \emptyset$ for each set $S \in \mathcal{S}$. We call H *minimal* if no strict subset is a hitting set. All minimal hitting sets can efficiently be computed with the algorithm HS-DAG (Greiner, Smith, Wilkerson 1989; Reiter 1987).

All MMSCDs but of $\emptyset$ are computable by the powering.

Proposition 8. $X^\mathcal{I} \equiv^\emptyset \exists^{\mathsf{sim}}(\wp(\mathcal{I}), X)$ *if* $\emptyset \neq X \subseteq \mathsf{Dom}(\mathcal{I})$.

Axiomatization of CIs by Means of FCA

The MMSCDs allow us to restrict attention to CIs of a particular form, viz. the set $\{ C \sqsubseteq C^{\mathcal{I}\mathcal{I}} \mid C \text{ is an } \mathcal{EL}_{\mathsf{si}} \text{ CD} \}$ would already be a CI base if it was finite. To see this, consider a CI $C \sqsubseteq D$ satisfied in $\mathcal{I}$. Then $C^{\mathcal{I}\mathcal{I}} \sqsubseteq^\emptyset D$ by (M2) and thus $C \sqsubseteq C^{\mathcal{I}\mathcal{I}}$ entails $C \sqsubseteq D$. The Galois property (G7) further ensures that every CI $C \sqsubseteq C^{\mathcal{I}\mathcal{I}}$ is satisfied in $\mathcal{I}$.

Each MMSCD $X^\mathcal{I}$ is either $\bot$ or, according to Proposition 8, a conjunction of CNs and existential restrictions $\exists r.Y^\mathcal{I}$. For this reason, we let $\mathbf{M}$ consist of $\bot$, all CNs, and all $\exists r.Y^\mathcal{I}$ where r is a RN and Y is a non-empty subset of $\mathsf{Dom}(\mathcal{I})$. This definition is up to equivalence, i.e. if $Y^\mathcal{I} \equiv^\emptyset Z^\mathcal{I}$ for two subsets $Y, Z \subseteq \mathsf{Dom}(\mathcal{I})$, then it suffices that $\mathbf{M}$ contains the attributes $\exists r.Y^\mathcal{I}$ for all RNs r.

Due to our choice of $\mathbf{M}$, we can now represent the conclusion $C^{\mathcal{I}\mathcal{I}}$ of any CI $C \sqsubseteq C^{\mathcal{I}\mathcal{I}}$ as a conjunction of atoms in $\mathbf{M}$, but this is not always possible for the premise C. We instead use the *partial closure* $C^{[\mathcal{I}\mathcal{I}]}$ which is closed everywhere above the root: $\bot^{[\mathcal{I}\mathcal{I}]} := \bot$ and, if $C \neq \bot$, then $C^{[\mathcal{I}\mathcal{I}]} := \bigsqcap\{ A \mid A \in \mathsf{Conj}(C) \} \sqcap \bigsqcap\{ \exists r.D^{\mathcal{I}\mathcal{I}} \mid \exists r.D \in \mathsf{Conj}(C) \}$.

Lemma 9. $\{ C^{[\mathcal{I}\mathcal{I}]} \sqsubseteq C^{\mathcal{I}\mathcal{I}} \mid C$ *is an* $\mathcal{EL}_{\mathsf{si}}$ *CD* $\}$ *is a CI base.*

Since the latter CI base consists of CIs between conjunctions over $\mathbf{M}$, Lemma 6 implies that we can obtain other, usually smaller CI bases by rewriting an implication base $\mathcal{B}$ of the induced context $\mathbb{K}_\mathcal{I}$ into the TBox $\bigsqcap \mathcal{B}$.

Moreover, we take the TBox $\mathcal{T}$ into account by transforming it into the set $\mathcal{L}_{\mathcal{I},\mathcal{T}}$ consisting of the implications

- $\mathsf{Conj}(C^{[\mathcal{I}\mathcal{I}]}) \to \{ E \mid E \in \mathbf{M} \text{ and } C \sqsubseteq^\mathcal{T} E \}$ for each CI $C \sqsubseteq D$ in $\mathcal{T}$
- $\{E\} \to \{F\}$ for each two $E, F \in \mathbf{M}$ with $E \sqsubseteq^\emptyset F$.

Theorem 10. *For each finite interpretation $\mathcal{I}$ and each $\mathcal{EL}^\bot_{\mathsf{si}}$ TBox $\mathcal{T}$ of which $\mathcal{I}$ is a model, the TBox $\mathsf{Can}(\mathcal{I}, \mathcal{T}) := \bigsqcap \mathsf{Can}(\mathbb{K}_\mathcal{I}, \mathcal{L}_{\mathcal{I},\mathcal{T}})$ is a CI base of $\mathcal{I}$ relative to $\mathcal{T}$. It is called* canonical CI base *and can be computed in time that is exponential in $\mathsf{Dom}(\mathcal{I})$ and polynomial in $\mathcal{T}$. If all CIs in $\mathcal{T}$ have the form $C \sqsubseteq D^{[\mathcal{I}\mathcal{I}]}$, then it contains the fewest CIs among all CI bases of $\mathcal{I}$ relative to $\mathcal{T}$. Furthermore, there are finite interpretations that have no polynomial-size CI base.*

Of course, we can strengthen the given TBox $\mathcal{T}$ by replacing every CI $C \sqsubseteq D$ with $C \sqsubseteq D^{[\mathcal{I}\mathcal{I}]}$ and then compute a minimal CI base of the interpretation $\mathcal{I}$ relative to this stronger TBox. Alternatively, we could compute the CI base relative to the unmodified TBox $\mathcal{T}$ and afterwards remove redundant CIs, which follow from others, but it is unclear whether this yields a CI base with the fewest possible number of CIs.

The next example shows that the computed CI base might not be minimal if not every CI in $\mathcal{T}$ has the form $C \sqsubseteq D^{[\mathcal{I}\mathcal{I}]}$.

Example 11. Consider the following interpretation $\mathcal{I}$.

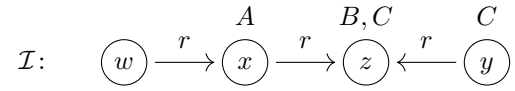

We further have the TBox $\mathcal{T} := \{A \sqsubseteq \exists r.B\}$ of which $\mathcal{I}$ is a model. Our goal is to compute the canonical CI base. We therefore first determine all MMSCDs, these are:

- $\{w\}^{\mathcal{I}} \equiv^{\emptyset} \exists r.(A \sqcap \exists r.(B \sqcap C))$
- $\{x\}^{\mathcal{I}} \equiv^{\emptyset} A \sqcap \exists r.(B \sqcap C)$
- $\{y\}^{\mathcal{I}} \equiv^{\emptyset} C \sqcap \exists r.(B \sqcap C)$
- $\{z\}^{\mathcal{I}} \equiv^{\emptyset} B \sqcap C$
- $\{x, y\}^{\mathcal{I}} \equiv^{\emptyset} \exists r.(B \sqcap C)$
- $\{y, z\}^{\mathcal{I}} \equiv^{\emptyset} C$
- $\{w, x, y\}^{\mathcal{I}} \equiv^{\emptyset} \exists r.\top$
- $\{w, x, y, z\}^{\mathcal{I}} \equiv^{\emptyset} \top$

We thus obtain the following induced context $\mathbb{K}_{\mathcal{I}}$.

$\mathbb{K}_{\mathcal{I}}$	$\top$	A	B	C	$\exists r.\{w\}^{\mathcal{I}}$	$\exists r.\{x\}^{\mathcal{I}}$	$\exists r.\{y\}^{\mathcal{I}}$	$\exists r.\{z\}^{\mathcal{I}}$	$\exists r.\{x,y\}^{\mathcal{I}}$	$\exists r.\{y,z\}^{\mathcal{I}}$	$\exists r.\{w,x,y\}^{\mathcal{I}}$	$\exists r.\{w,x,y,z\}^{\mathcal{I}}$
w	·	·	·	·	·	×	·	·	×	·	×	×
x	·	×	·	·	·	·	·	×	·	×	·	×
y	·	·	·	×	·	·	·	×	·	×	·	×
z	·	·	×	×	·	·	·	·	·	·	·	·

The implication set $\mathcal{L}_{\mathcal{I},\mathcal{T}}$ consists of all $\{E\} \to \{F\}$ where $E \sqsubseteq^{\emptyset} F$ and of $\{A\} \to \{\exists r.\{w, x, y, z\}^{\mathcal{I}}\}$. Note that the latter evaluates to $\{A\} \to \{\exists r.\top\}$ and thus does not fully capture the CI $A \sqsubseteq \exists r.B$. By transforming the canonical implication base of $\mathbb{K}_{\mathcal{I}}$ relative to $\mathcal{L}_{\mathcal{I},\mathcal{T}}$, we obtain the canonical CI base of $\mathcal{I}$ relative to $\mathcal{T}$ with the following CIs:

- $B \sqsubseteq C$
- $A \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \exists r.\{z\}^{\mathcal{I}} \sqcap \exists r.\{y, z\}^{\mathcal{I}}$
- $C \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \exists r.\{z\}^{\mathcal{I}} \sqcap \exists r.\{y, z\}^{\mathcal{I}}$
- $\exists r.\{y, z\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \exists r.\{z\}^{\mathcal{I}}$
- $\exists r.\{w, x, y\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \exists r.\{x\}^{\mathcal{I}} \sqcap \exists r.\{x, y\}^{\mathcal{I}}$
- $A \sqcap C \sqcap \exists r.\{z\}^{\mathcal{I}} \sqcap \exists r.\{y, z\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \bot$
- $B \sqcap C \sqcap \exists r.\{z\}^{\mathcal{I}} \sqcap \exists r.\{y, z\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \bot$
- $\exists r.\{w\}^{\mathcal{I}} \sqcap \exists r.\{x\}^{\mathcal{I}} \sqcap \exists r.\{x, y\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \bot$
- $\exists r.\{x\}^{\mathcal{I}} \sqcap \exists r.\{z\}^{\mathcal{I}} \sqcap \exists r.\{x, y\}^{\mathcal{I}} \sqcap \exists r.\{y, z\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y\}^{\mathcal{I}} \sqcap \exists r.\{w, x, y, z\}^{\mathcal{I}} \sqsubseteq \bot$

Specifically the second CI is superfluous as it can be deduced from the others. To see this, first note that it can be simplified to $A \sqcap \exists r.\top \sqsubseteq \exists r.(B \sqcap C)$. Now, the premise $A \sqcap \exists r.\top$ is subsumed by $\exists r.B$ (since $\mathcal{T}$ contains $A \sqsubseteq \exists r.B$) and thus subsumed by $\exists r.(B \sqcap C)$ (since the CI base contains $B \sqsubseteq C$), which is the conclusion of the second CI.

To obtain a minimal CI base, we could replace the conclusion of the CI $A \sqsubseteq \exists r.B$ in $\mathcal{T}$ with $(\exists r.B)^{[\mathcal{I}\mathcal{I}]} = \exists r.(B \sqcap C)$. Then, the implication set $\mathcal{L}_{\mathcal{I},\mathcal{T}}$ would contain $\{A\} \to \{\exists r.(B \sqcap C)\}$ in place of $\{A\} \to \{\exists r.\top\}$ and so the replaced CI could be fully captured.

The canonical CI base contains $\mathcal{EL}^{\bot}_{\mathsf{si}}$ CIs. We rewrite it into $\mathcal{EL}$ to gain support by state-of-the-art reasoners, such as ELK. We subdivide the rewriting into two steps. First, we replace the premises in $\mathsf{Can}(\mathcal{I}, \mathcal{T})$ by suitable $\mathcal{EL}$ CDs. Minimality is preserved since such a replacement does not change the number of CIs in the base. In particular, we can replace, in every conjunction $\bigsqcap \mathbf{C}$ that occurs as a premise, each $\exists r.X^{\mathcal{I}}$ by $\exists r.(X^{\mathcal{I}}\restriction_n)$ with the following choices for n.

- n is minimal such that $(X^{\mathcal{I}}\restriction_n)^{\mathcal{I}} = X^{\mathcal{I}\mathcal{I}}$, which can be determined by trying non-negative integers in ascending order and picking the first for which the MMSCD $X^{\mathcal{I}}$ and its unfolding $X^{\mathcal{I}}\restriction_n$ have the same extension in $\mathcal{I}$.
- $n := 2^{|\mathsf{Dom}(\mathcal{I})|} \cdot |\mathsf{Dom}(\mathcal{I})| + 1$ (Baader, Distel 2008)
- n is obtained from the MVF measure based on lengths of simple paths in $\mathcal{I}$ or powers of $\mathcal{I}$, seen as graphs (Guimarães, Ozaki, Persia, Sertkaya 2021).

Next, we replicate the cyclic structures within the conclusions by means of $\mathcal{EL}$ CIs. Assume that $\mathcal{B}$ is obtained from $\mathsf{Can}(\mathcal{I}, \mathcal{T})$ by replacing the premises as above. We obtain an $\mathcal{EL}$ TBox $\mathcal{B}'$ with auxiliary CNs that entails the same CIs as follows. For each CI $C \sqsubseteq \bot$ in $\mathcal{B}$, we add $C \sqsubseteq \bot$ to $\mathcal{B}'$. For each CI $C \sqsubseteq \bigsqcap \mathbf{D}$ in $\mathcal{B}$, we add the following CIs to $\mathcal{B}'$:

- $C \sqsubseteq \bigsqcap\{\, A \mid A \in \mathbf{D} \,\} \sqcap \bigsqcap\{\, \exists r.X \mid \exists r.X^{\mathcal{I}} \in \mathbf{D} \,\}$
- $Y \sqsubseteq \bigsqcap\{\, A \mid Y \in A^{\wp(\mathcal{I})} \,\} \sqcap \bigsqcap\{\, \exists r.Z \mid (Y, Z) \in r^{\wp(\mathcal{I})} \,\}$ for all Y reachable from any X in $\wp(\mathcal{I})$ with $\exists r.X^{\mathcal{I}} \in \mathbf{D}$

Due to *structural sharing* this transformation is considerably smaller than the known one (Baader, Distel 2008).

Axiomatization of RRs and RIs

As $\mathcal{I}$ satisfies a RR $\top \sqsubseteq \forall r.C$ if $\bigcup\{\, r^{\mathcal{I}}(x) \mid x \in \mathsf{Dom}(\mathcal{I}) \,\} \subseteq C^{\mathcal{I}}$, the most specific RR on r uses the MMSCD $Y^{\mathcal{I}}$ of $Y := \bigcup\{\, r^{\mathcal{I}}(x) \mid x \in \mathsf{Dom}(\mathcal{I}) \,\}$ in place of C. We thus add these RRs to the CI base, possibly after replacing $\top \sqsubseteq \forall r.Y^{\mathcal{I}}$ with $\top \sqsubseteq \forall r.Y$ and the CIs describing all Z reachable from Y.

Next, we show how RIs can be completely axiomatized. Since $\mathcal{I}$ is finite, all RIs are regular in the following sense.

Proposition 12. *For every RN s, the language $L_{\mathcal{I}}(s) := \{\, r_1 \cdots r_n \mid \mathcal{I}$ satisfies the RI $r_1 \circ \cdots \circ r_n \sqsubseteq s \,\}$ is regular and is accepted by a finite automaton $\mathfrak{A}_s$ of exponential size.*

The automata for all RNs can be converted into a complete set of RIs, which we add to the CI base. We therefore use the automata states as auxiliary RNs, and introduce the RI $p \circ r \sqsubseteq q$ for each transition (p, r, q), the RI $\varepsilon \sqsubseteq i$ for each initial state i, and the RI $f \sqsubseteq s$ for each final state f in $\mathfrak{A}_s$. It remains open how these RIs can be rewritten into equivalent RIs without auxiliary RNs, but we believe this is possible. However, many reasoners transform the RIs in a given ontology into automata anyway and for these the above RIs are advantageous since the automata can easily be read off.

Theorem 13. *For each finite interpretation $\mathcal{I}$, a complete TBox of $\mathcal{EL}$ CIs, RRs, and RIs satisfied in $\mathcal{I}$ can be computed in exponential time. There are finite interpretations for which such a TBox cannot be of polynomial size.*

Implementation Details

Several steps employ maximal simulations and therefore a performant algorithm for computing these is advantageous. We adapt an approach to computing simulations between graphs (M. R. Henzinger, T. A. Henzinger, Kopke 1995)

such that it works with interpretations (which can be seen as labelled graphs) and runs in parallel on multiple threads.

Computing the canonical CI base needs exponential time in the worst case. We reduce the input interpretation $\mathcal{I}$ to save computation time. The key observation is that we can group together all objects in $\mathsf{Dom}(\mathcal{I})$ satisfying the same CDs. By doing so, no counterexamples against CIs satisfied in $\mathcal{I}$ are removed, and also no new counterexamples against satisfied CIs are introduced. However, instead of checking infinitely many CIs the following characterization comes to the rescue: if $\mathcal{I}$ and $\mathcal{J}$ are finite, then $(\mathcal{I}, x) \preceq (\mathcal{J}, y)$ iff. $x \in C^{\mathcal{I}}$ implies $y \in C^{\mathcal{J}}$ for all C (Lutz, Wolter 2010). Thus, in order to decide whether two objects x and y in $\mathcal{I}$ satisfy the same CDs, we check if the maximal simulation $\mathfrak{S}_{\mathcal{I}}$ on $\mathcal{I}$ contains (x, y) as well as (y, x). Based on this we compute a *weak reduction* of $\mathcal{I}$ in polynomial time and show that it fulfils a certain minimality condition.

To compute the attribute set $\mathbf{M}$, we should not naïvely go through all subsets of $\mathsf{Dom}(\mathcal{I})$ and compute their MMSCDs because there are exponentially many subsets and MMSCDs of different subsets are often equivalent. The following consequence of Properties (G1)–(G7) helps us.

Lemma 14. *The mapping* $\phi_{\mathcal{I}} \colon X \mapsto X^{\mathcal{I}\mathcal{I}}$ *is a closure operator on* $\mathsf{Dom}(\mathcal{I})$.

In order to avoid computing duplicates, we enumerate all closures of $\phi_{\mathcal{I}}$ with a multi-threaded version of the algorithm Fast Close-by-One (FCbO) (Krajča, Outrata, Vychodil 2010). By Property (G4) and Proposition 8 all MMSCDs are then obtained from the closures $X^{\mathcal{I}\mathcal{I}}$ as $\exists^{\mathsf{sim}}(\wp(\mathcal{I}), X^{\mathcal{I}\mathcal{I}})$. The operator $\phi_{\mathcal{I}}$ is computed with the maximal simulation $\mathfrak{S}_{\wp(\mathcal{I}),\mathcal{I}}$ from the powering $\wp(\mathcal{I})$ to $\mathcal{I}$.

Lemma 15. $X^{\mathcal{I}\mathcal{I}} = \mathfrak{S}_{\wp(\mathcal{I}),\mathcal{I}}(X)$ *if* $\emptyset \neq X \subseteq \mathsf{Dom}(\mathcal{I})$.

To avoid fully constructing the exponentially-large powering, we lazily build only the part reachable from X when a closure $X^{\mathcal{I}\mathcal{I}}$ is computed.

For some datasets even these lazily constructed parts of $\wp(\mathcal{I})$ are so large that they cannot be computed within reasonable time limits. In order to detect such cases beforehand and to not waste computation time, the prototype approximates, for the current object set at which the powering is to be expanded, the number of successors — if it is larger than 10,000,000, the computation will be aborted. The resulting CI base will, however, not be complete anymore since some attributes required in the set $\mathbf{M}$ could not be computed. But, if completeness comes for the price of extremely large CIs, which might not have practical relevance or suffer from overfitting, then one can probably dispense with this goal. Moreover, we also allowed to manually specify a smaller limit on the number of successors and thereby to further restrict the size of CIs in the base. We turned this bound into a *conjunction size limit* by also counting the CNs that label the particular object set in $\wp(\mathcal{I})$. It remains unclear to which extent completeness is lost, and we leave the investigation as future research. We expect that completeness is still guaranteed for all CIs that obey the conjunction size limit, but modifications to the method might be needed to achieve this.

To compute the canonical implication base of the induced context $\mathbb{K}_{\mathcal{I}}$ we employ the algorithm LinCbO (Janoštík, Konečný, Krajča 2021a,b, 2022). It is based on Close-by-One (CbO) (Kuznetsov 1993) and closures w.r.t. implications are computed with an improved version of LinClosure that reuses counters. Like all CbO-based algorithms, it uses the *canonicity test* to avoid duplicate computations of the same closure. This test is integrated into the modified LinClosure sub-routine, which enables *early stop* of unnecessary computation branches, and is additionally supported by *pruning rules*. However, we need to extend LinCbO with support for background implications.

Variations

Some CIs in the canonical CI base $\mathsf{Can}(\mathcal{I}, \mathcal{T})$ are disjointness axioms $C \sqsubseteq \bot$, which express that no objects in $\mathcal{I}$ are described by C. Sometimes only the other CIs $C \sqsubseteq D$ are desired as they describe the implications between CDs that are satisfied and also witnessed in $\mathcal{I}$. We have seen in experiments that more than half of the computation time is required for generating disjointness axioms. It is cheaper to compute only the witnessed CIs since some intermediate computation steps can be stopped early.

Definition 16. A CI $C \sqsubseteq D$ is *witnessed* in $\mathcal{I}$ if $C^{\mathcal{I}} \neq \emptyset$ and $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. A TBox is *witnessed complete* for $\mathcal{I}$ if it entails all CIs that are witnessed in $\mathcal{I}$. A *witnessed CI base* of $\mathcal{I}$ relative to $\mathcal{T}$ is a TBox $\mathcal{B}$ that consists of witnessed CIs and for which $\mathcal{B} \cup \mathcal{T}$ is witnessed complete.

Proposition 17. *Let* $\mathsf{Can}_+(\mathcal{I}, \mathcal{T})$ *consist of all witnessed CIs in the canonical CI base* $\mathsf{Can}(\mathcal{I}, \mathcal{T})$. *Then* $\mathsf{Can}_+(\mathcal{I}, \mathcal{T})$ *is a witnessed CI base of* $\mathcal{I}$ *relative to* $\mathcal{T}$. *Among all witnessed CI bases of* $\mathcal{I}$ *relative to* $\mathcal{T}$ *it contains the fewest CIs.*

To compute $\mathsf{Can}_+(\mathcal{I}, \mathcal{T})$, we should exclude from the attribute set $\mathbf{M}$ all CDs not describing any object in $\mathcal{I}$. We can avoid the expensive computation of such attributes $\exists r.X^{\mathcal{I}}$ by a modification of the employed algorithm FCbO. Furthermore, since the other algorithm LinCbO enumerates the implications in a sub-order of set inclusion $\subseteq$ on the premises, we can stop a computation branch as soon as an implication has been found that is not witnessed in the induced context $\mathbb{K}_{\mathcal{I}}$ (as it would yield a CI that is not witnessed in $\mathcal{I}$). In a similar way, we incorporate a conjunction size limit ℓ, viz. we stop a branch when encountering an implication with more than ℓ attributes in the premise.

If, instead, we want a CI base that is still complete for disjointness axioms but which need not be minimal, then we compute the attribute set $\mathbf{M}$ as usual, but before building the induced context $\mathbb{K}_{\mathcal{I}}$ from it we remove every CD C not satisfied in $\mathcal{I}$, except $\bot$, and store the *fast disjointness axiom* $C \sqsubseteq \bot$ in an intermediate set that we will later add to the computed CI base. Since thereby the size of $\mathbb{K}_{\mathcal{I}}$ is often significantly reduced, the computation of the canonical implication base is much faster. The downside is, however, that the final CI base is larger.

Apart from bounding the conjunction size, another effective way to avoid the axiomatization of impractically huge CIs is to limit the role depth. Specifically, the *role depth* of an $\mathcal{EL}$ CD is the maximal number of nestings of existential restrictions. By modifications to the approach on Page 4, we can also compute a CI base w.r.t. a role-depth bound $n \geq 0$

which is, however, only guaranteed to be complete for all CIs bounded by n. The case without a known TBox $\mathcal{T}$ has already been considered (Borchmann, Distel, Kriegel 2016). We show how such an existing TBox can be taken into account, yielding a minimal CI base as for the unrestricted case if $\mathcal{T}$ also satisfies the role-depth bound n.

Theorem 18. *Given a finite interpretation $\mathcal{I}$, an $\mathcal{EL}^{\bot}_{\mathsf{si}}$ TBox $\mathcal{T}$ of which $\mathcal{I}$ is a model, and a number $n \geq 0$, then a CI base of $\mathcal{I}$ relative to $\mathcal{T}$ for role depth n can be computed in time that is exponential in $\mathsf{Dom}(\mathcal{I})$ and polynomial in $\mathcal{T}$ and n. If all CIs in $\mathcal{T}$ are of the form $C \sqsubseteq D^{[\mathcal{I}\mathcal{I}]}\restriction_n$ and bounded by n, then it contains the fewest CIs among all CI bases of $\mathcal{I}$ relative to $\mathcal{T}$ for n. Furthermore, there are finite interpretations of which no CI base for n has polynomial size.*

As an application of Theorem 18, we can keep the CIs in a base small by iteratively axiomatizing CIs from a given interpretation $\mathcal{I}$. We therefore increase the role-depth bound in each step (starting with 0) and take all CIs in $\mathcal{T}$ as well as the CIs from all previous steps as background knowledge. This guarantees a CI base that is complete for all CIs when the role-depth bound $2^{|\mathsf{Dom}(\mathcal{I})|} \cdot |\mathsf{Dom}(\mathcal{I})| + 1$ has been reached (Baader, Distel 2008). Alternatively, we could stop earlier and as last step compute the canonical CI base from Theorem 10 relative to $\mathcal{T}$ and all CIs from the previous steps.

Experimental Evaluation

We implemented (Kriegel 2023b) the axiomatization method in the programming language Scala 3 and we evaluate the prototype with the plethora of ontologies from real-world applications used in the ORE 2015 Reasoner Competition (Parsia, Matentzoglu, Gonçalves, Glimm, Steigmiller 2017). This collection is split into OWL 2 EL and OWL 2 DL ontologies. The former cannot contain any CIs not expressible in $\mathcal{EL}$. For the latter, we syntactically transform as many axioms as possible into $\mathcal{EL}$ and remove the others. There is no best way to do this since optimal finite $\mathcal{EL}$ approximations need not exist (Haga, Lutz, Marti, Wolter 2020). Removal of unsupported axioms makes these ontologies weaker in the sense that some logical consequences are lost; however, no new, undesired consequences are thereby introduced. Each test dataset is derived from such an ontology, viz. we treat the ABox as interpretation $\mathcal{I}$ (under closed-world assumption) and the TBox $\mathcal{T}$ as existing knowledge. If necessary, we saturate $\mathcal{I}$ by means of the CIs in $\mathcal{T}$. Altogether we obtain 614 test datasets with up to 747,998 objects, of which 446 (72.64 %) are acyclic. The average number of triples per object varies from slightly over 0 up to 25.39.

The prototype supports three modes in which disjointness axioms are not computed (None), computed in the fast way (Fast), or computed in the canonical way (Can.) as per Theorem 10. It further allows for specifying a role depth bound and a conjunction size limit. During the experiment, we used all three modes and the settings $(0, 32)$, $(1, 8)$, $(1, 32)$, $(2, 32)$, $(\infty, 32)$, (∞, ∞) where the first parameter is the role depth bound and the second is the conjunction size limit. For every dataset, the prototype was executed once for each configuration (mode and parameters).

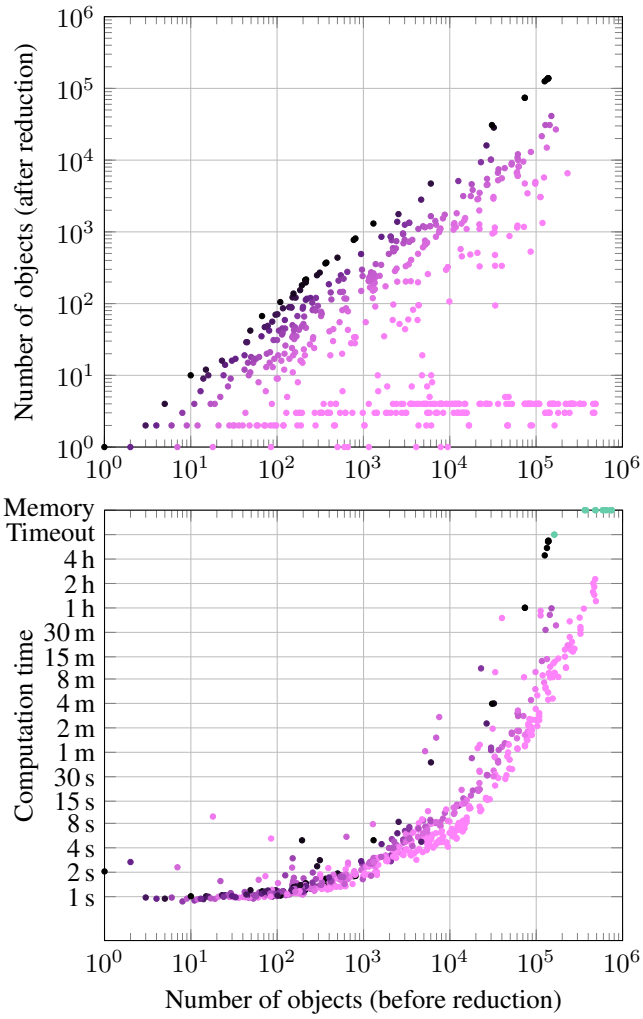

Figure 1: Computing reductions of the test datasets

The experiments were run on a small, old computer server with two Intel Xeon E5-2640 processors (each with 6 CPU cores, hyper-threading, 15 MB cache, 2.80 GHz frequency, boost up to 3.00 GHz) and 96 GB DDR3-SDRAM main memory. Modern laptops have faster processors but usually only a smaller amount of main memory. As runtime environment we used Oracle GraalVM EE 22.3.0 (Java 19.0.1).

For 599 (97.56 %) of the 614 test datasets the prototype successfully computed the (weak) reduction. Figure 1 shows computation times as well as size changes. In many cases, the number of objects was significantly reduced and often by more than one order of magnitude. Several reductions contain fewer than ten objects, meaning that there is only a small variety of different types of objects. We ignored these for the subsequent experiment steps. Reductions could not be computed for 13 (2.12 %) larger datasets with more than 300,000 objects due to out-of-memory errors (limit: 80 GB), and for 2 (0.33 %) datasets due to timeouts (limit: 8 hours).

Figure 2 shows the computation times for the CI bases (without reduction) and their sizes, including failures due

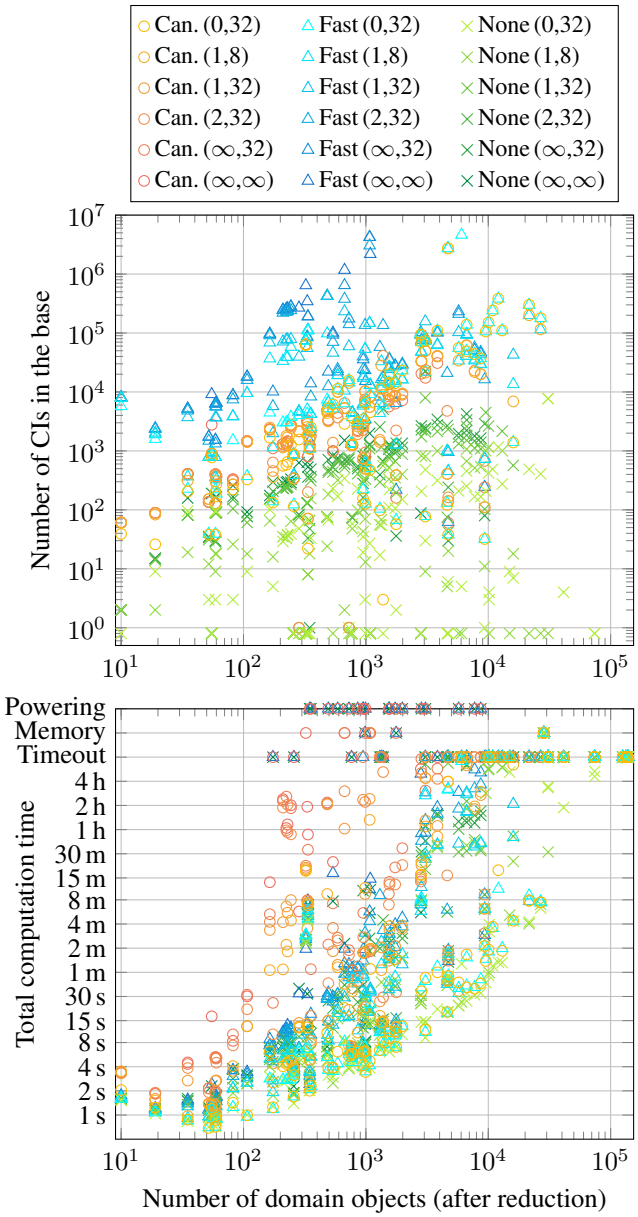

Figure 2: Computing CI bases of the reduced test datasets

to timeouts (limit: 8 hours), out-of-memory errors (limit: 80 GB), or powering-too-large exceptions (conjunction size limit: 10,000,000). However, we did not implement and measure the rewriting of the CI base into $\mathcal{EL}$, nor the axiomatization of RRs and RIs. Computation finished for all reduced datasets with no more than 100 objects. For reduced datasets with up to 1,000 objects, the first errors due to insufficient computing resources occurred without a role-depth bound. Between 1,000 and 10,000 objects, computations failed without restrictions, but otherwise succeeded in the majority of cases. Reduced datasets with more than 10,000 objects could only sometimes be axiomatized with very restricted settings, given 8 hours time and 80 GB memory.

In summary, we clearly see that computation resources can be saved if no disjointness axioms are wanted, or if they are computed in the fast way. Furthermore, the parameters allow us to control the overall resource consumption on the one hand, but also the size and number of the CIs in the final base on the other hand. We can avoid the computation of huge CIs that might not have any practical relevance.

More details concerning all intermediate computation steps can be found in the extended version (Kriegel 2023a); a brief summary is as follows. Most expensive are the computation of all MMSCDs and of the can. implication base of the induced context. In more restricted settings the CDs that could differentiate the objects are smaller and fewer, resulting in fewer MMSCDs that can be computed in less time, also yielding smaller induced contexts. Given all MMSCDs, computing the induced context and the background implications is cheap and needs only seconds for most datasets.

Future Prospects

That the theoretical approach itself can be extended to more expressive DLs has already been proven, but it is unclear whether such an extended approach can still be efficiently implemented and used in practice. From the perspective of this article, this seems possible for DLs characterized by simulations, e.g. $\mathcal{ELI}$ or Horn-$\mathcal{ALC}$.

Regarding the presented approach, an interesting question for future research would be whether one can give any kind of completeness guarantee if a conjunction size limit is used (e.g. every CI that also satisfies the limit is entailed). A smaller task can be to investigate how range restrictions and role inclusions can be integrated into the background knowledge after they have been computed but prior to axiomatizing the CIs, preferably yielding an overall minimal base.

It should be investigated whether the canonical CI base can be obtained more efficiently from the fast CI base by means of the algorithm in (Rudolph 2007). The witnessed CI base ignores all disjointness axioms. One could restrict it even more by requiring that the number of objects satisfying a CI premise must exceed an absolute or relative limit.

Furthermore, the computation can be speed-up with even faster FCA algorithms for enumerating closures. The employed LinCbO algorithm is currently the fastest algorithm for computing the canonical implication base, but it is unfortunately only single-threaded. Developing a multi-threaded variant is thus another future goal. It might already help to change its depth-first behaviour. Apart from that one could use a faster programming language (like C++), more computation time, a faster server, or optimize the prototype.

A CI $C \sqsubseteq D$ is *confident* if the ratio $|(C \sqcap D)^{\mathcal{I}}|/|C^{\mathcal{I}}|$ exceeds a pre-defined limit but need not be 100 % (Borchmann 2013). Since a confident CI base extends a canonical CI base by CIs of the form $X^{\mathcal{I}} \sqsubseteq Y^{\mathcal{I}}$, the prototype could be upgraded as it already computes all MMSCDs $X^{\mathcal{I}}$ and $Y^{\mathcal{I}}$.

We have not considered keys supported by the OWL 2 EL profile. Learning of keys from RDF data using FCA has been addressed in (Abbas, Bazin, David, Napoli 2021, 2022; Atencia, David, Euzenat, Napoli, Vizzini 2020). To apply this approach to DL and OWL it must be extended towards complex DL concepts in place of RDF classes.

Acknowledgements

The author thanks the anonymous reviewers for their helpful feedback. The author is funded by Deutsche Forschungsgemeinschaft (DFG) in Projects 430150274 (Repairing Description Logic Ontologies) and 389792660 (TRR 248: Foundations of Perspicuous Software Systems).

References

Abbas, N.; Bazin, A.; David, J.; Napoli, A. 2021. Non-Redundant Link Keys in RDF Data: Preliminary Steps. In *Proc. of FCA4AI*, 125–130. https://ceur-ws.org/Vol-2972/paper12.pdf.

Abbas, N.; Bazin, A.; David, J.; Napoli, A. 2022. A Study of the Discovery and Redundancy of Link Keys Between Two RDF Datasets Based on Partition Pattern Structures. In *Proc. of CLA*, 175–189. https://ceur-ws.org/Vol-3308/Paper14.pdf.

Abboud, R.; Ceylan, İ. İ.; Lukasiewicz, T.; Salvatori, T. 2020. BoxE: A Box Embedding Model for Knowledge Base Completion. In *Proc. of NeurIPS*.

Alghamdi, G.; Schmidt, R. A.; Del-Pinto, W.; Gao, Y. 2021. Upwardly Abstracted Definition-Based Subontologies. In *Proc. of K-CAP*, 209–216. https://doi.org/10.1145/3460210.3493564.

Asaadi, S.; Giesbrecht, E.; Rudolph, S. 2023. Compositional matrix-space models of language: Definitions, properties, and learning methods. *Nat. Lang. Eng.*, 29(1): 32–80. https://doi.org/10.1017/S1351324921000206.

Atencia, M.; David, J.; Euzenat, J.; Napoli, A.; Vizzini, J. 2020. Link key candidate extraction with relational concept analysis. *Discret. Appl. Math.*, 273: 2–20. https://doi.org/10.1016/j.dam.2019.02.012.

Baader, F. 1995. Computing a Minimal Representation of the Subsumption Lattice of all Conjunctions of Concepts Defined in a Terminology. In *Proc. of KRUSE*, 168–178.

Baader, F. 2003. Terminological Cycles in a Description Logic with Existential Restrictions. In *Proc. of IJCAI*, 325–330. http://ijcai.org/Proceedings/03/Papers/048.pdf.

Baader, F.; Brandt, S.; Lutz, C. 2005. Pushing the $\mathcal{EL}$ Envelope. In *Proc. of IJCAI*, 364–369. http://ijcai.org/Proceedings/05/Papers/0372.pdf.

Baader, F.; Brandt, S.; Lutz, C. 2008. Pushing the $\mathcal{EL}$ Envelope Further. In *Proc. of OWLED*. http://ceur-ws.org/Vol-496/owled2008dc_paper_3.pdf.

Baader, F.; Distel, F. 2008. A Finite Basis for the Set of $\mathcal{EL}$-Implications Holding in a Finite Model. In *Proc. of ICFCA*, 46–61. https://doi.org/10.1007/978-3-540-78137-0_4.

Baader, F.; Distel, F. 2009. Exploring Finite Models in the Description Logic $\mathcal{EL}_{\text{gfp}}$. In *Proc. of ICFCA*, 146–161. https://doi.org/10.1007/978-3-642-01815-2_12.

Baader, F.; Ganter, B.; Sertkaya, B.; Sattler, U. 2007. Completing Description Logic Knowledge Bases Using Formal Concept Analysis. In *Proc. of IJCAI*, 230–235. http://ijcai.org/Proceedings/07/Papers/035.pdf.

Baader, F.; Horrocks, I.; Lutz, C.; Sattler, U. 2017. *An Introduction to Description Logic*. https://doi.org/10.1017/9781139025355.

Baader, F.; Molitor, R. 2000. Building and Structuring Description Logic Knowledge Bases Using Least Common Subsumers and Concept Analysis. In *Proc. of ICCS*, 292–305. https://doi.org/10.1007/10722280_20.

Baader, F.; Sertkaya, B. 2004. Applying Formal Concept Analysis to Description Logics. In *Proc. of ICFCA*, 261–286. https://doi.org/10.1007/978-3-540-24651-0_24.

Beeri, C.; Bernstein, P. A. 1979. Computational Problems Related to the Design of Normal Form Relational Schemas. *ACM Trans. Database Syst.*, 4(1): 30–59. https://doi.org/10.1145/320064.320066.

Bogaerts, B.; Jakubowski, M.; Van den Bussche, J. 2022. SHACL: A Description Logic in Disguise. In *Proc. of LPNMR*, 75–88. https://doi.org/10.1007/978-3-031-15707-3_7.

Borchmann, D. 2013. Towards an Error-Tolerant Construction of $\mathcal{EL}^{\perp}$-Ontologies from Data Using Formal Concept Analysis. In *Proc. of ICFCA*, 60–75. https://doi.org/10.1007/978-3-642-38317-5_4.

Borchmann, D. 2015. Exploring Faulty Data. In *Proc. of ICFCA*, 219–235. https://doi.org/10.1007/978-3-319-19545-2_14.

Borchmann, D.; Distel, F.; Kriegel, F. 2016. Axiomatisation of general concept inclusions from finite interpretations. *J. Appl. Non Class. Logics*, 26(1): 1–46. https://doi.org/10.1080/11663081.2016.1168230.

Chen, J.; Jiménez-Ruiz, E.; Horrocks, I.; Chen, X.; Myklebust, E. B. 2023. An assertion and alignment correction framework for large scale knowledge bases. *Semantic Web*, 14(1): 29–53. https://doi.org/10.3233/SW-210448.

Cuenca Grau, B.; Horrocks, I.; Kazakov, Y.; Sattler, U. 2008. Modular Reuse of Ontologies: Theory and Practice. *J. Artif. Intell. Res.*, 31: 273–318. https://doi.org/10.1613/jair.2375.

Distel, F. 2011. *Learning description logic knowledge bases from data using methods from formal concept analysis*. Doctoral thesis. Technische Universität Dresden. https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa-70199.

Funk, M.; Jung, J. C.; Lutz, C. 2022. Frontiers and Exact Learning of ELI Queries under DL-Lite Ontologies. In *Proc. of IJCAI*, 2627–2633. https://doi.org/10.24963/ijcai.2022/364.

Funk, M.; Jung, J. C.; Lutz, C.; Pulcini, H.; Wolter, F. 2019. Learning Description Logic Concepts: When can Positive and Negative Examples be Separated? In *Proc. of IJCAI*, 1682–1688. https://doi.org/10.24963/ijcai.2019/233.

Ganter, B. 1984. *Two Basic Algorithms in Concept Analysis.* FB4-Preprint 831. Technische Hochschule Darmstadt. https://doi.org/10.1007/978-3-642-11928-6_22.

Ganter, B.; Wille, R. 1999. *Formal Concept Analysis – Mathematical Foundations*. https://doi.org/10.1007/978-3-642-59830-2.

Greiner, R.; Smith, B. A.; Wilkerson, R. W. 1989. A Correction to the Algorithm in Reiter's Theory of Diagnosis. *Artif. Intell.*, 41(1): 79–88. https://doi.org/10.1016/0004-3702(89)90079-9.

Guigues, J.-L.; Duquenne, V. 1986. Famille minimale d'implications informatives résultant d'un tableau de données binaires. *Mathématiques et Sciences Humaines*, 95: 5–18. http://www.numdam.org/item/MSH_1986__95__5_0.pdf.

Guimarães, R.; Ozaki, A.; Persia, C.; Sertkaya, B. 2021. Mining $\mathcal{EL}$ Bases with Adaptable Role Depth. In *Proc. of AAAI*, 6367–6374. https://doi.org/10.1609/aaai.v35i7.16790.

Haga, A.; Lutz, C.; Marti, J.; Wolter, F. 2020. A Journey into Ontology Approximation: From Non-Horn to Horn. In *Proc. of IJCAI*, 1827–1833. https://doi.org/10.24963/ijcai.2020/253.

Henzinger, M. R.; Henzinger, T. A.; Kopke, P. W. 1995. Computing Simulations on Finite and Infinite Graphs. In *Proc. of FOCS*, 453–462. https://doi.org/10.1109/SFCS.1995.492576.

Hitzler, P.; Krötzsch, M.; Rudolph, S. 2010. *Foundations of Semantic Web Technologies*. http://www.semantic-web-book.org/.

Horrocks, I.; Olivares, J.; Cocchi, V.; Motik, B.; Roy, D. 2022. The Dow Jones Knowledge Graph. In *Proc. of ESWC*, 427–443. https://doi.org/10.1007/978-3-031-06981-9_25.

Jackermeier, M.; Chen, J.; Horrocks, I. 2023. Box^{2}EL: Concept and Role Box Embeddings for the Description Logic $\mathcal{EL}^{++}$. https://doi.org/10.48550/arXiv.2301.11118.

Jain, N.; Kalo, J.; Balke, W.; Krestel, R. 2021. Do Embeddings Actually Capture Knowledge Graph Semantics? In *Proc. of ESWC*, 143–159. https://doi.org/10.1007/978-3-030-77385-4_9.

Janoštík, R.; Konečný, J.; Krajča, P. 2021a. LinCbO: Fast algorithm for computation of the Duquenne-Guigues basis. *Inf. Sci.*, 572: 223–240. https://doi.org/10.1016/j.ins.2021.04.104.

Janoštík, R.; Konečný, J.; Krajča, P. 2021b. Pruning Techniques in LinCbO for Computation of the Duquenne-Guigues Basis. In *Proc. of ICFCA*, 91–106. https://doi.org/10.1007/978-3-030-77867-5_6.

Janoštík, R.; Konečný, J.; Krajča, P. 2022. Pruning techniques in LinCbO for the computation of the Duquenne-Guigues basis. *Inf. Sci.*, 616: 182–203. https://doi.org/10.1016/j.ins.2022.10.057.

Jimeno-Yepes, A.; Jiménez-Ruiz, E.; Llavori, R. B.; Rebholz-Schuhmann, D. 2009. Reuse of terminological resources for efficient ontological engineering in Life Sciences. *BMC Bioinform.*, 10(S-10): 4. https://doi.org/10.1186/1471-2105-10-S10-S4.

Kazakov, Y.; Klinov, P. 2015. Advancing ELK: Not Only Performance Matters. In *Proc. of DL*. http://ceur-ws.org/Vol-1350/paper-27.pdf.

Kazakov, Y.; Krötzsch, M.; Simančik, F. 2014. The Incredible ELK – From Polynomial Procedures to Efficient Reasoning with $\mathcal{EL}$ Ontologies. *J. Autom. Reason.*, 53(1): 1–61. https://doi.org/10.1007/s10817-013-9296-3.

Klarman, S.; Britz, K. 2015. Towards Unsupervised Ontology Learning from Data. In *Proc. of DARe*. https://ceur-ws.org/Vol-1423/DARe-15_5.pdf.

Konev, B.; Lutz, C.; Ozaki, A.; Wolter, F. 2017. Exact Learning of Lightweight Description Logic Ontologies. *J. Mach. Learn. Res.*, 18: 201:1–201:63. http://jmlr.org/papers/v18/16-256.html.

Krajča, P.; Outrata, J.; Vychodil, V. 2010. Advances in Algorithms Based on CbO. In *Proc. of CLA*, 325–337. http://ceur-ws.org/Vol-672/paper29.pdf.

Kriegel, F. 2017. Acquisition of Terminological Knowledge from Social Networks in Description Logic. *Formal Concept Analysis of Social Networks*, 97–142. https://doi.org/10.1007/978-3-319-64167-6_5.

Kriegel, F. 2019. Joining Implications in Formal Contexts and Inductive Learning in a Horn Description Logic. In *Proc. of ICFCA*, 110–129. https://doi.org/10.1007/978-3-030-21462-3_9.

Kriegel, F. 2023a. *Efficient Axiomatization of OWL 2 EL Ontologies from Data by means of Formal Concept Analysis (Extended Version)*. LTCS-Report 23-01. Technische Universität Dresden. https://doi.org/10.25368/2023.214.

Kriegel, F. 2023b. *Efficient Axiomatization of OWL 2 EL Ontologies from Data by means of Formal Concept Analysis (Prototypical Implementation)*. https://github.com/francesco-kriegel/efficient-axiomatization-of-owl2el-ontologies-from-data.

Krötzsch, M. 2019. Too Much Information: Can AI Cope with Modern Knowledge Graphs? In *Proc. of ICFCA*, 17–31. https://doi.org/10.1007/978-3-030-21462-3_2.

Kuznetsov, S. O. 1993. A fast algorithm for computing all intersections of objects from an arbitrary semilattice. (1): 17–20.

Lutz, C.; Piro, R.; Wolter, F. 2010. Enriching $\mathcal{EL}$-Concepts with Greatest Fixpoints. In *Proc. of ECAI*, 41–46. https://doi.org/10.3233/978-1-60750-606-5-41.

Lutz, C.; Wolter, F. 2010. Deciding inseparability and conservative extensions in the description logic $\mathcal{EL}$. *J. Symb. Comput.*, 45(2): 194–228. https://doi.org/10.1016/j.jsc.2008.10.007.

Lutz, C.; Wolter, F. 2011. Foundations for Uniform Interpolation and Forgetting in Expressive Description Logics. In *Proc. of IJCAI*, 989–995. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-170.

Parsia, B.; Matentzoglu, N.; Gonçalves, R. S.; Glimm, B.; Steigmiller, A. 2017. The OWL Reasoner Evaluation (ORE) 2015 Competition Report. *J. Autom. Reason.*, 59(4): 455–482. https://doi.org/10.1007/s10817-017-9406-8. The test ontologies are available from https://doi.org/10.5281/zenodo.18578.

Reiter, R. 1987. A Theory of Diagnosis from First Principles. *Artif. Intell.*, 32(1): 57–95. https://doi.org/10.1016/0004-3702(87)90062-2. See the erratum (Greiner, Smith, Wilkerson 1989).

Rudolph, S. 2004. Exploring Relational Structures Via $\mathcal{FLE}$. In *Proc. of ICCS*, 196–212. https://doi.org/10.1007/978-3-540-27769-9_13.

Rudolph, S. 2007. Some Notes on Pseudo-closed Sets. In *Proc. of ICFCA*, 151–165. https://doi.org/10.1007/978-3-540-70901-5_10.

Sarker, M. K.; Krisnadhi, A. A.; Hitzler, P. 2016. OWLAx: A Protege Plugin to Support Ontology Axiomatization through Diagramming. In *Proc. of ISWC Demonstrations Track*. https://ceur-ws.org/Vol-1690/paper83.pdf.

Shi, J.; Chen, J.; Dong, H.; Khan, I.; Liang, L.; Zhou, Q.; Wu, Z.; Horrocks, I. 2023. Subsumption Prediction for E-Commerce Taxonomies. In *Proc. of ESWC*, 244–261. https://doi.org/10.1007/978-3-031-33455-9_15.

Stumme, G. 1996. Attribute Exploration with Background Implications and Exceptions, 457–469. https://doi.org/10.1007/978-3-642-80098-6_39.

Wild, M. 1994. A Theory of Finite Closure Spaces Based on Implications. *Advances in Mathematics*, 108(1): 118–139. https://doi.org/10.1006/aima.1994.1069.

Wille, R. 1982. Restructuring Lattice Theory: An Approach Based on Hierarchies of Concepts. In *Ordered Sets*, 445–470. https://doi.org/10.1007/978-3-642-01815-2_23.

Zarrieß, B.; Turhan, A. 2013. Most Specific Generalizations w.r.t. General $\mathcal{EL}$-TBoxes. In *Proc. of IJCAI*, 1191–1197. http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6709.

Zhao, Y.; Schmidt, R. A.; Wang, Y.; Zhang, X.; Feng, H. 2020. A Practical Approach to Forgetting in Description Logics with Nominals. In *Proc. of AAAI*, 3073–3079. https://ojs.aaai.org/index.php/AAAI/article/view/5702.