

Tractable Diversity: Scalable Multiperspective Ontology Management via Standpoint \mathcal{EL}

Lucía Gómez Álvarez and Sebastian Rudolph and Hannes Strass

Computational Logic Group, TU Dresden, Germany

{lucia.gomez_alvarez, sebastian.rudolph, hannes.strass}@tu-dresden.de

Abstract

The tractability of the lightweight description logic \mathcal{EL} has allowed for the construction of large and widely used ontologies that support semantic interoperability. However, comprehensive domains with a broad user base are often at odds with strong axiomatisations otherwise useful for inferencing, since these are usually context dependent and subject to diverging perspectives.

In this paper we introduce *Standpoint \mathcal{EL}* , a multi-modal extension of \mathcal{EL} that allows for the integrated representation of domain knowledge relative to diverse, possibly conflicting *standpoints* (or contexts), which can be hierarchically organised and put in relation to each other. We establish that *Standpoint \mathcal{EL}* still exhibits \mathcal{EL} 's favourable PTIME standard reasoning, whereas introducing additional features like empty standpoints, rigid roles, and nominals makes standard reasoning tasks intractable.

1 Introduction

In many subfields of artificial intelligence, ontologies are used to provide a formal representation of a shared vocabulary, give meaning to its terms, and describe the relations between them. To this end, one of the most prominent and successful class of logic-based knowledge representation formalisms are *description logics* (DLs) [Baader *et al.*, 2017; Rudolph, 2011], which provide the formal basis for the most recent version of the Web Ontology Language OWL 2 [OWL Working Group, 2009].

Among the most prominent families of DLs used today is \mathcal{EL} [Baader *et al.*, 2005], which is the formal basis of OWL 2 EL [Motik *et al.*, 2009a], a popular tractable profile of OWL 2. One of the main appeals of \mathcal{EL} is that basic reasoning tasks can be performed in polynomial time with respect to the size of the ontology, enabling reasoning-supported creation and maintenance of very large ontologies. An example of this is the healthcare ontology SNOMED CT [Donnelly, 2006], with worldwide adoption and a broad user base comprising clinicians, patients, and researchers among others.

However, when modelling comprehensive ontologies like SNOMED CT, one is usually facing issues related to context

or perspective-dependent knowledge as well as ambiguity of language [Schulz *et al.*, 2017]. For instance, the concept Tumour might denote a process or a piece of tissue; Allergy may denote an allergic reaction or just an allergic disposition.

In a similar vein, the decentralised nature of the Semantic Web has led to the generation of various ontologies of overlapping knowledge that inevitably reflect different points of view. For instance, an initiative has attempted to integrate the FMA1140 (Foundational Model of Anatomy), SNOMED-CT, and the NCI (National Cancer Institute Thesaurus) into a single combined version called LargeBio and reported ensuing challenges [Osman *et al.*, 2021]. In this context, frameworks supporting the integrated representation of multiple perspectives seem preferable to recording the distinct views in a detached way, but also to entirely merging them at the risk of causing inconsistencies or unintended consequences.

To this end, Gómez Álvarez and Rudolph [2021] proposed *standpoint logic*, a formalism inspired by the theory of supervaluationism [Fine, 1975] and rooted in modal logic, which allows for the simultaneous representation of multiple, potentially contradictory, viewpoints in a unified way and the establishment of alignments between them. This is achieved by extending the base language with labelled modal operators, where propositions $\Box_S\phi$ and $\Diamond_S\phi$ express information relative to the *standpoint* S and read, respectively: “according to S , it is *unequivocal/conceivable* that ϕ ”. Semantically, standpoints are represented by sets of *precisifications*,¹ such that $\Box_S\phi$ and $\Diamond_S\phi$ hold if ϕ is true in all/some of the precisifications associated with S . Consider the following example.

Example 1 (Tumour Disambiguation). *Two derivatives of the SNOMED CT ontology (SN) model tumours differently. According to TP, a Tumour is a process by which abnormal or damaged cells grow and multiply (1), yet according to TT, a Tumour is a lump of tissue (2).*

$$\Box_{TP}[\text{Tumour} \sqsubseteq \text{AbnormalGrowthProcess}] \quad (1)$$

$$\Box_{TT}[\text{Tumour} \sqsubseteq \text{Tissue}] \quad (2)$$

Both interpretations inherit the axioms of the original SNOMED CT (3) and are such that: if according to SN something is arguably both a Tumour and a Tissue, then it (unequivocally) is a Tumour according to TT (4). The respective

¹Precisifications are analogous to the *worlds* of modal-logic frameworks with possible-worlds semantics.

assertion is made for TP (5). But Tissue and Process are disjoint categories according to SN (6).

$$(TP \preceq SN) \quad (TT \preceq SN) \quad (3)$$

$$\diamond_{SN}[\text{Tumour} \sqcap \text{PhysicalObject}] \sqsubseteq \square_{TT}[\text{Tumour}] \quad (4)$$

$$\diamond_{SN}[\text{Tumour} \sqcap \text{Process}] \sqsubseteq \square_{TP}[\text{Tumour}] \quad (5)$$

$$\square_{SN}[\text{Tissue} \sqcap \text{Process}] \sqsubseteq \perp \quad (6)$$

While clearly incompatible, both perspectives are semantically close and we can establish relations between them. For instance, we might assert that something is unequivocally the product of a Tumour (process) according to TP if and only if it is arguably a Tumour (tissue) according to TT (7). Also, we may specify a subsumption between the classes of unequivocal instances of Tissue according to TT and to TP (8).

$$\square_{TP}[\exists \text{ProductOf} . \text{Tumour}] \equiv \diamond_{TT}[\text{Tumour}] \quad (7)$$

$$\square_{TT}[\text{Tissue}] \sqsubseteq \square_{TP}[\text{Tissue}] \quad (8)$$

When recording clinical findings, clinicians may use ambiguous language, so an automated knowledge extraction service may obtain the following from text and annotated scans:

$$\square_{SN}\{\text{Patient}(p), \text{HasPart}(p, a), \text{Colon}(a)\} \quad (9)$$

$$\diamond_{SN}\{\text{HasPart}(a, b), \text{Tumour}(b), \text{PhysicalObject}(b)\} \quad (10)$$

The logical statements (1)–(10), which formalise Example 1 by means of a standpoint-enhanced \mathcal{EL} description logic, are not inconsistent, so all axioms can be jointly represented. Let us now illustrate the use of standpoint logic for reasoning with and across individual perspectives.

Example 2 (Continued from Example 1). In this case, we can disambiguate the information given by Axiom (10) using Axiom (3) and Axiom (4), which entail that according to TT, b is unequivocally a tumour, $\square_{TT}\text{Tumour}(b)$, and with Axiom (2) also a tissue, $\square_{TT}\text{Tissue}(b)$. Moreover, we can use the “bridges” to switch to another perspective. From Axiom (8), it is clear that according to TP, b is also a tissue, $\square_{TP}\text{Tissue}(b)$, and from Axiom (7) b is the product of a tumour, $\square_{TP}\exists \text{ProductOf} . \text{Tumour}(b)$. Then Axiom (1) yields

$$\square_{TP}\exists \text{ProductOf} . (\text{Tumour} \sqcap \text{AbnormalGrowthProcess})(b).$$

The statement $\square_{SN}[\text{Tumour} \sqcap \text{Process}](d)$, in contrast, will trigger an inconsistency thanks to Axiom (6), which prevents the evaluation of Tumour simultaneously as a Tissue and a Process and Axiom (2), which states that according to some interpretations, a Tumour is a Tissue. Finally, a user (e.g. a specific clinic, CL) may inherit the SNOMED CT ($\text{CL} \preceq \text{SN}$) and establish further axioms, e.g.

$$\square_{CL}[\text{Patient} \sqcap \exists \text{HasPart} . (\text{Colon} \sqcap \diamond_{SN}\exists \text{HasPart} . \text{Tumour}) \\ \sqsubseteq \exists \text{AssociatedWith} . \text{ColonCancerRisk}],$$

to identify patients with cancer risk. Here, Axiom (9) lets us infer that $\square_{CL}\exists \text{AssociatedWith} . \text{ColonCancerRisk}(p)$. \diamond

The need of handling multiple perspectives in the Semantic Web has led to several (non-modal) logic-based proposals. The closest regarding goals are multi-viewpoint ontologies [Hemam and Boufaïda, 2011; Hemam, 2018], which model the intuition of viewpoints in a tailored extension of OWL for which no complexity bounds are given. Similar problems

are also addressed in the more extensive work on contextuality (e.g. C-OWL and Distributed ontologies [Bouquet *et al.*, 2003; Borgida and Serafini, 2003] and the Contextualised Knowledge Repository (CKR) [Serafini and Homola, 2012]). These frameworks focus on contextual and distributed reasoning and range between different levels of expressivity for modelling the structure of contexts and the bridges between them. In the context of scalable reasoning, one should highlight the implementations that provide support for OWL2-RL based CKR defeasible reasoning [Bozzato *et al.*, 2018].

As for modal logics, their suitability to model perspectives and contexts in a natural way is obvious [Klarman and Gutiérrez-Basulto, 2013; Gómez Álvarez and Rudolph, 2021], they are well-known in the community and their semantics is well-understood. Yet, the interplay between DL constructs and modalities is often not well-behaved and can easily endanger the decidability of reasoning tasks or increase their complexity [Baader and Ohlbach, 1995; Mosurović, 1999; Wolter and Zakharyashev, 1999]. Notable examples are NEXPTIME-completeness of the multi-modal description logic $\mathbf{K}_{\mathcal{ALC}}$ [Lutz *et al.*, 2002] and 2EXPTIME-completeness of $\mathcal{ALC}_{\mathcal{ALC}}$ [Klarman and Gutiérrez-Basulto, 2013], a modal contextual logic framework in the style proposed by McCarthy [McCarthy and Buvac, 1998].

In this work, we focus on the framework of *standpoint logics* [Gómez Álvarez and Rudolph, 2021], which are modal logics, too, but come with a simplified Kripke semantics. Recently, Gómez Álvarez *et al.* [2022] introduced *First-Order Standpoint Logic* (FOSL) and showed favourable complexity results for its *sentential* fragments,² which disallow modal operators being applied to formulas with free variables. In particular, adding sentential standpoints does not increase the complexity for fragments that are NP-hard. Yet, a fine-grained terminological alignment between different perspectives requires concepts preceded by modal operators, as in Axiom (7), leading to non-sentential fragments of FOSL.

Our paper is structured as follows. After introducing the syntax and semantics of Standpoint \mathcal{EL} ($\mathcal{S}_{\mathcal{EL}}$) and a suitable normal form (Section 2), we establish our main result: satisfiability checking in $\mathcal{S}_{\mathcal{EL}}$ is PTIME-complete. We show this by providing a worst-case optimal tableau-based algorithm (Section 3) that takes inspiration from the *quasi-model* based methods [Wolter and Zakharyashev, 1998] as used for $\mathbf{K}_{\mathcal{ALC}}$ [Lutz *et al.*, 2002], but differs in its specifics. Our approach builds a *quasi-model* from a graph of (*quasi*) *domain elements*, which are annotated with various constraints, to then reconstruct the worlds or, in our case, precisifications. We also show that introducing additional features such as empty standpoints, rigid roles, and nominals make standard reasoning tasks intractable (Section 4). In Section 5, we conclude the paper with a discussion of future work, including efficient approaches for reasoner implementations. Altogether, this paper provides a clear pathway for making scalable multiperspective ontology management possible.

An extended version of the paper with proofs of all results is available at <https://arxiv.org/abs/2302.13187>.

²This includes the sentential standpoint variant of the expressive DL SRCTQ_b s, a logical basis of OWL 2 DL [Motik *et al.*, 2009b].

2 Syntax, Semantics, and Normalisation

We now introduce syntax and semantics of Standpoint \mathcal{EL} (referred to as $\mathbb{S}_{\mathcal{EL}}$) and propose a normal form that is useful for subsequent algorithmic considerations.

Syntax

A *Standpoint DL vocabulary* is a traditional DL vocabulary consisting of sets N_C of *concept names*, N_R of *role names*, and N_I of *individual names*, extended it by an additional set N_S of *standpoint names* with $*$ in N_S . A *standpoint operator* is of the form \diamond_s (“diamond”) or \square_s (“box”) with $s \in N_S$; we use \odot_s to refer to either. A *concept term* is defined via

$$C ::= \top \mid \perp \mid A \mid C_1 \sqcap C_2 \mid \exists R.C \mid \odot_s[C]$$

where $A \in N_C$ and $R \in N_R$. A *general concept inclusion (GCI)* is of the form $\odot_s[C \sqsubseteq D]$, where C and D are concept terms.³ A *concept assertion* is of the form $\odot_s[C(a)]$ while a *role assertion* is of the form $\odot_s[R(a, b)]$, where $a, b \in N_I$, C is a concept term, and $R \in N_R$. A *sharpening statement* is of the form $s \preceq s'$ where $s, s' \in N_S$.

A $\mathbb{S}_{\mathcal{EL}}$ *knowledge base* is a tuple $\mathcal{K} = \langle \mathcal{S}, \mathcal{T}, \mathcal{A} \rangle$, where \mathcal{T} is a set of GCIs, called *TBox*; \mathcal{A} is a set of (concept or role) assertions, called *ABox*; and \mathcal{S} is a set of sharpening statements, called *SBox*. We refer to arbitrary statements from \mathcal{K} as *axioms*. Since the axiom types in \mathcal{S} , \mathcal{T} , and \mathcal{A} are syntactically well-distinguished, we sometimes identify \mathcal{K} as $\mathcal{S} \cup \mathcal{T} \cup \mathcal{A}$. Note that all axioms except sharpening statements are preceded by modal operators (“modalised” for short). In case the preceding operator happens to be \square_* , we may omit it.

Semantics

The semantics of $\mathbb{S}_{\mathcal{EL}}$ is defined via standpoint structures. For a Standpoint DL vocabulary $\langle N_C, N_R, N_I, N_S \rangle$, a *description logic standpoint structure* is a tuple $\mathfrak{D} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ where:

- Δ is a non-empty set, the *domain* of \mathfrak{D} ;
- Π is a set, called the *precisifications* of \mathfrak{D} ;
- σ is a function mapping each standpoint name to a non-empty subset of Π ;⁴
- γ is a function mapping each precisification from Π to an “ordinary” DL interpretation $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ over the domain Δ , where the interpretation function $\cdot^{\mathcal{I}}$ maps:
 - each concept name $A \in N_C$ to a set $A^{\mathcal{I}} \subseteq \Delta$,
 - each role name $R \in N_R$ to a binary relation $R^{\mathcal{I}} \subseteq \Delta \times \Delta$,
 - each individual name $a \in N_I$ to an element $a^{\mathcal{I}} \in \Delta$,
and we require $a^{\gamma(\pi)} = a^{\gamma(\pi')}$ for all $\pi, \pi' \in \Pi$ and $a \in N_I$.

Note that by this definition, individual names (also referred to as constants) are interpreted rigidly, i.e., each individual name a is assigned the same $a^{\gamma(\pi)} \in \Delta$ across all precisifications $\pi \in \Pi$. We will refer to this uniform $a^{\gamma(\pi)}$ by $a^{\mathfrak{D}}$.

For each $\pi \in \Pi$, the interpretation function $\cdot^{\mathcal{I}}$ for $\mathcal{I} = \gamma(\pi)$ is extended to concept terms via structural induction as follows:

³The square brackets $[\dots]$ indicate the scope of the modality, as the same modalities may be used inside concept terms.

⁴As shown in Section 4, allowing for “empty standpoints” immediately incurs intractability, even for an otherwise empty vocabulary.

$$\begin{aligned} \top^{\mathcal{I}} &:= \Delta & (\diamond_s C)^{\mathcal{I}} &:= \bigcup_{\pi' \in \sigma(s)} C^{\gamma(\pi')} \\ \perp^{\mathcal{I}} &:= \emptyset & (\square_s C)^{\mathcal{I}} &:= \bigcap_{\pi' \in \sigma(s)} C^{\gamma(\pi')} \\ (C_1 \sqcap C_2)^{\mathcal{I}} &:= C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}} \\ (\exists R.C)^{\mathcal{I}} &:= \{ \delta \in \Delta \mid \langle \delta, \varepsilon \rangle \in R^{\mathcal{I}} \text{ for some } \varepsilon \in C^{\mathcal{I}} \} \end{aligned}$$

We observe that modalised concepts $\odot_s C$ are interpreted uniformly across all precisifications $\pi \in \Pi$, which allows us to denote their extensions with $(\odot_s C)^{\mathfrak{D}}$.

A DL standpoint structure \mathfrak{D} *satisfies a sharpening statement* $s \preceq s'$, written as $\mathfrak{D} \models s \preceq s'$, iff $\sigma(s) \subseteq \sigma(s')$. For the other axiom types, satisfaction by \mathfrak{D} is defined as follows:

$$\begin{aligned} \mathfrak{D} \models \square_s[C \sqsubseteq D] &:\iff C^{\gamma(\pi)} \subseteq D^{\gamma(\pi)} \text{ for each } \pi \in \sigma(s) \\ \mathfrak{D} \models \diamond_s[C \sqsubseteq D] &:\iff C^{\gamma(\pi)} \subseteq D^{\gamma(\pi)} \text{ for some } \pi \in \sigma(s) \\ \mathfrak{D} \models \square_s[C(a)] &:\iff a^{\mathfrak{D}} \in \bigcap_{\pi \in \sigma(s)} C^{\gamma(\pi)} (= (\square_s C)^{\mathfrak{D}}) \\ \mathfrak{D} \models \diamond_s[C(a)] &:\iff a^{\mathfrak{D}} \in \bigcup_{\pi \in \sigma(s)} C^{\gamma(\pi)} (= (\diamond_s C)^{\mathfrak{D}}) \\ \mathfrak{D} \models \square_s[R(a, b)] &:\iff \langle a^{\mathfrak{D}}, b^{\mathfrak{D}} \rangle \in \bigcap_{\pi \in \sigma(s)} R^{\gamma(\pi)} \\ \mathfrak{D} \models \diamond_s[R(a, b)] &:\iff \langle a^{\mathfrak{D}}, b^{\mathfrak{D}} \rangle \in \bigcup_{\pi \in \sigma(s)} R^{\gamma(\pi)} \end{aligned}$$

As usual, \mathfrak{D} is a *model of \mathcal{S}* iff it satisfies every sharpening statement in \mathcal{S} ; it is a *model of \mathcal{T}* iff it satisfies every GCI $\tau \in \mathcal{T}$; it is a *model of \mathcal{A}* iff it satisfies every assertion $\alpha \in \mathcal{A}$; it is a *model of $\mathcal{K} = \langle \mathcal{S}, \mathcal{T}, \mathcal{A} \rangle$* (written $\mathfrak{D} \models \mathcal{K}$) iff it is a model of \mathcal{S} and a model of \mathcal{T} and a model of \mathcal{A} . Figure 1 illustrates a model of Example 1.




Figure 1: *Model of Example 1.* The curly brackets represent σ , mapping standpoint names to subsets of Π . The dotted lines represent the DL interpretation at each precisification, with concepts in blue and roles denoted by arrows. An individual name mapping is indicated under the element line.

Our investigations regarding reasoning in $\mathbb{S}_{\mathcal{EL}}$ will focus on standpoint versions of the well-known standard reasoning tasks, and we will make use of variations of established techniques to (directly or indirectly) reduce all of them to the first.

Knowledge base satisfiability: *Given a knowledge base \mathcal{K} , is there a DL standpoint structure \mathfrak{D} such that $\mathfrak{D} \models \mathcal{K}$?*

Axiom entailment: *Given \mathcal{K} and some SBox, TBox, or ABox axiom ϕ , does $\mathcal{K} \models \phi$ hold, that is, is it the case that for every model \mathfrak{D} of \mathcal{K} we have $\mathfrak{D} \models \phi$?*

To show that axiom entailment can be polynomially reduced to knowledge base unsatisfiability, we exhibit for every axiom type ϕ a knowledge base $\mathcal{K}_{-\phi}$ such that $\mathcal{K} \models \phi$ coincides with unsatisfiability of $\mathcal{K} \cup \mathcal{K}_{-\phi}$.

$$\begin{aligned} \mathcal{K}_{s \preceq s'} &:= \{ \diamond_s[\top \sqsubseteq \tilde{A}], \square_{s'}[\tilde{A} \sqsubseteq \perp] \} \\ \mathcal{K}_{\odot_s[C \sqsubseteq D]} &:= \{ \tilde{A} \sqsubseteq C, \tilde{A} \sqcap D \sqsubseteq \perp, \odot_s^d[\top \sqsubseteq \exists \tilde{R}.\tilde{A}] \} \\ \mathcal{K}_{\odot_s[C(a)]} &:= \{ \tilde{A} \sqcap C \sqsubseteq \perp, \odot_s^d[\tilde{A}(a)] \} \\ \mathcal{K}_{\odot_s[R(a, b)]} &:= \{ \tilde{B}(b), \tilde{A} \sqcap \exists R.\tilde{B} \sqsubseteq \perp, \odot_s^d[\tilde{A}(a)] \} \end{aligned}$$

Therein, $\Box_s^d := \Diamond_s$ and $\Diamond_s^d := \Box_s$ are modality duals, and \tilde{A}, \tilde{B} denote fresh concept names and \tilde{R} a fresh role name.

Concept satisfiability (w.r.t. \mathcal{K}): Given \mathcal{K} and a modalised concept term C , is there a model \mathcal{D} of \mathcal{K} with $C^{\mathcal{D}} \neq \emptyset$?

This task can be solved by checking the axiom entailment $\mathcal{K} \models \Box_*[C \sqsubseteq \perp]$. If the entailment holds, then C is unsatisfiable w.r.t. \mathcal{K} , otherwise it is satisfiable.

Instance retrieval: Given \mathcal{K} and a modalised concept term C , obtain all $a \in N_I$ with $a^{\mathcal{D}} \in C^{\mathcal{D}}$ for every model \mathcal{D} of \mathcal{K} . This task can be solved by checking, for all individuals a , if the entailment $\mathcal{K} \models \Box_*[C(a)]$ holds and returning all such a .

Normalisation

Before we can describe a PTIME algorithm for checking satisfiability of $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ KBs, we define an appropriate normal form.

Definition 1 (Normal Form of $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ Knowledge Bases). A TBox \mathcal{T} is in normal form iff, for all its GCIs $\Box_s[C \sqsubseteq D]$,

1. C is of the form A , $\exists R.A$ or $A \sqcap A'$ with $A, A' \in N_C \cup \{\top\}$,
2. D is of the form B , $\exists R.B$, $\Diamond_{s'}B$ or $\Box_{s'}B$ with $B \in N_C \cup \{\perp\}$, and
3. at least one of C, D is from $N_C \cup \{\top, \perp\}$;

where $R \in N_R$, and $s, s' \in N_S$.

An ABox \mathcal{A} is in normal form iff all assertions have the form $\Box_s[A(a)]$ or $\Box_s[R(a, b)]$ for $a, b \in N_I$, $A \in N_C$, and $R \in N_R$. $\mathcal{K} = \langle \mathcal{S}, \mathcal{T}, \mathcal{A} \rangle$ is in normal form whenever \mathcal{T} and \mathcal{A} are. \diamond

For a given $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ knowledge base $\mathcal{K} = \langle \mathcal{S}, \mathcal{T}, \mathcal{A} \rangle$, we can compute its normal form by exhaustively applying the following transformation rules (where “rule application” means that the axiom on the left-hand side is replaced with the set of axioms on the right-hand side),

$$\Diamond_s[C(a)] \rightarrow \{v \preceq s, \Box_v[C(a)]\} \quad (11)$$

$$\Diamond_s[R(a, b)] \rightarrow \{v \preceq s, \Box_v[R(a, b)]\} \quad (12)$$

$$\Diamond_s[C \sqsubseteq D] \rightarrow \{v \preceq s, \Box_v[C \sqsubseteq D]\} \quad (13)$$

$$\Box_s[\tilde{C}(a)] \rightarrow \{\Box_s[A(a)], \Box_s[A \sqsubseteq \tilde{C}]\} \quad (14)$$

$$\Box_s[B \sqsubseteq \exists R.\tilde{C}] \rightarrow \{\Box_s[B \sqsubseteq \exists R.A], \Box_s[A \sqsubseteq \tilde{C}]\} \quad (15)$$

$$\Box_s[B \sqsubseteq C \sqcap D] \rightarrow \{\Box_s[B \sqsubseteq A], \Box_s[A \sqsubseteq C], \Box_s[A \sqsubseteq D]\} \quad (16)$$

$$\Box_s[C \sqsubseteq \odot_u \tilde{D}] \rightarrow \{\Box_s[C \sqsubseteq \odot_u A], \Box_s[A \sqsubseteq \tilde{D}]\} \quad (17)$$

$$\Box_s[C \sqsubseteq \top] \rightarrow \emptyset \quad \text{and} \quad \Box_s[\perp \sqsubseteq D] \rightarrow \emptyset \quad (18)$$

$$\Box_s[\exists R.\tilde{C} \sqsubseteq D] \rightarrow \{\Box_s[\tilde{C} \sqsubseteq A], \Box_s[\exists R.A \sqsubseteq D]\} \quad (19)$$

$$\Box_s[\tilde{C} \sqcap D \sqsubseteq E] \rightarrow \{\Box_s[\tilde{C} \sqsubseteq A], \Box_s[A \sqcap D \sqsubseteq E]\} \quad (20)$$

$$\Box_s[\Diamond_u C \sqsubseteq D] \rightarrow \{\Box_u[C \sqsubseteq \Box_* A], \Box_s[A \sqsubseteq D]\} \quad (21)$$

$$\Box_s[\Box_u C \sqsubseteq D] \rightarrow \{v_0 \preceq u, v_1 \preceq u, \Box_u[C \sqsubseteq A], \Box_s[\Diamond_{v_0} A \sqcap \Diamond_{v_1} A \sqsubseteq D]\} \quad (22)$$

Therein, \tilde{C} and \tilde{D} stand for complex concept terms not contained in $N_C \cup \{\top\}$, whereas each occurrence of A on a right-hand side denotes the introduction of a fresh concept name; likewise, v , v_0 , and v_1 denote the introduction of a fresh standpoint name. Rule (20) is applied modulo commutativity of \sqcap . Most of the transformation rules should be intuitive (for the first three, keep in mind that standpoints must be nonempty). A notable exception is Rule (22), which

is crucial to remove boxes occurring with negative polarity. It draws some high-level inspiration from existing work on non-vacuous left-hand-side universal quantifiers in Horn DLs [Carral *et al.*, 2014], yet the argument for its correctness requires a much more intricate model-theoretic construction and crucially hinges on “Hornness” of \mathcal{K} and nonemptiness of standpoints. A careful analysis yields that the transformation has the desired semantic and computational properties.

Lemma 1. Every $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ knowledge base \mathcal{K} can be transformed into a $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ knowledge base \mathcal{K}' in normal form such that

- \mathcal{K}' is a $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ -conservative extension of \mathcal{K} ,
- the size of \mathcal{K}' is at most linear in the size of \mathcal{K} , and
- the transformation can be computed in PTIME.

While \mathcal{K}' being a $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ -conservative extension of \mathcal{K} brings about various valuable properties, what matters for our purposes is that this implies equisatisfiability of \mathcal{K} and \mathcal{K}' , thus we will not go into details about conservative extensions.

3 A Tableau Algorithm for Standpoint $\mathcal{E}\mathcal{L}$

We present a PTIME tableau decision algorithm for $\mathbb{S}_{\mathcal{E}\mathcal{L}}$. Complexity-optimal tableau algorithms have been proposed for description logics with modal operators applied to concepts and axioms such as $\mathbf{K}_{\mathcal{A}\mathcal{L}\mathcal{C}}$ [Lutz *et al.*, 2002], which is known to be in NEXPTIME. Our case cannot be treated in the same way, as we need to take greater care to show tractability in the end. Lutz *et al.* [2002] build a “quasi-model” from a tree of “quasi-worlds”, which is not as easily applicable in our case, so we follow a dual approach: we will build a *quasi-model* from a completion graph of (*quasi*) domain elements, where each of the latter is associated to a constraint system with assembled information regarding one individual’s specifics in each precisification. We begin with some definitions. Given a $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ knowledge base \mathcal{K} , denote by

- $\text{ST}_{\mathcal{K}}$ the elements of N_S occurring in \mathcal{K} together with $*$,
- $\text{IN}_{\mathcal{K}}$ the set of all individual names occurring in \mathcal{K} ,
- $\text{BC}_{\mathcal{K}}$ (*basic concepts*) the concept names used in \mathcal{K} , plus \top ,
- $\text{C}_{\mathcal{K}}$ (*set of concept terms*) the set of concept terms used in \mathcal{K} (with $\text{BC}_{\mathcal{K}} \subseteq \text{C}_{\mathcal{K}}$),
- $\text{SF}_{\mathcal{K}}$ the set of *subformulas* of \mathcal{K} , consisting of all axioms of \mathcal{K} with and without their outer standpoint modality.

A *constraint* for \mathcal{K} is of the form $(x : C)$, $(x : a)$, $(x : \phi)$, or $(x : s)$,⁵ where x is a variable, $C \in \text{C}_{\mathcal{K}}$ a concept, $a \in \text{IN}_{\mathcal{K}}$ an individual, $\phi \in \text{SF}_{\mathcal{K}}$ a formula, and $s \in \text{ST}_{\mathcal{K}}$ a standpoint name. *Constraint systems* are finite sets of constraints.

Definition 2 ((Initial) Constraint System for \mathcal{K}). The initial constraint system for \mathcal{K} , called $S_0^{\mathcal{K}}$, is the set

$$\{x_s : *, x_s : \top, x_s : \phi, x_s : s \mid \phi \in \mathcal{K}, s \in \text{ST}_{\mathcal{K}}\}$$

A constraint system for \mathcal{K} is a finite set S of constraints for \mathcal{K} such that $S_0^{\mathcal{K}} \subseteq S$ and $\{x : *, x : \top\} \subseteq S$ for each x in S . For a variable x , let $\text{st}_S(x) = \{s \mid (x : s) \in S\}$ be the standpoint signature of x in S . \diamond

Intuitively, each constraint system S produced by the algorithm corresponds to a domain element $\varepsilon \in \Delta$ and each variable x in S corresponds to some precisification π .

⁵For better legibility, we will sometimes omit the parentheses.




Figure 2: For a $\mathbb{S}_{\mathcal{EL}}$ knowledge base $\mathcal{K} = \{\top \sqsubseteq \diamond_s C, C \sqsubseteq \exists R.D\}$: completion graph \mathbf{CG} produced by the algorithm (left); quasi-model \mathcal{Q} over \mathbf{CG} with the runs $\Gamma = \{r_1, \dots, r_4\}$ represented by dotted lines (center); and model \mathcal{D} obtained by the construction in the proof of Theorem 2 (right). Constraints of the form $(x : \phi)$ in \mathbf{CG} and \mathcal{Q} are omitted for better readability.

Moreover, each constraint $x : X$ in S encodes information of ε in π . Namely, X may be an axiom that holds in π , a standpoint that contains π , or a concept expression of which ε is an instance in π . Initialising one variable per standpoint in the initial constraint system guarantees non-empty standpoints.

A constraint system is *complete* iff it satisfies every local completion rule from Figure 3. Local completion rules operate on constraint systems, while global rules involve more than one constraint system and operate on completion graphs.

Definition 3 (Completion Graph). An element label is a set L of triples of the form (C, s, x) , where $C \in \text{BC}_{\mathcal{K}}$ is a concept, $s \subseteq \text{ST}_{\mathcal{K}}$ is a set of standpoints, and x is a variable.

A quasi-role for a set Δ is a tuple $\langle \varepsilon, v, \varepsilon', v', R \rangle$ where v and v' are variables, $\varepsilon, \varepsilon' \in \Delta$, and R is a role name in \mathcal{K} .

A completion graph for \mathcal{K} is a tuple $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$, with Δ a non-empty set of elements; \mathcal{S} a map from Δ into constraint systems; \mathcal{L} a map from Δ into element labels; and \mathcal{R} a set of quasi-roles such that

- for all $\langle \varepsilon, v, \varepsilon', v', R \rangle \in \mathcal{R}$, $(v : s) \in \mathcal{S}(\varepsilon)$ iff $(v' : s) \in \mathcal{S}(\varepsilon')$;
- if $(C, s, x) \in \mathcal{L}(\varepsilon)$, then $\{x : C\} \cup \{x : s \mid s \in s\} \subseteq \mathcal{S}(\varepsilon)$. \diamond

For convenience of presentation, we use the shortcut $\text{st}_{\varepsilon}(v)$ for $\text{st}_{\mathcal{S}(\varepsilon)}(v)$ and for any $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$, we will refer to all $\varepsilon \in \Delta$ simply as elements of \mathbf{CG} .

\mathbf{CG} is said to be *locally complete* iff for every element ε in \mathbf{CG} , $\mathcal{S}(\varepsilon)$ is complete, and we call \mathbf{CG} *globally complete* iff it is locally complete and no global completion rule (see Figure 3) is applicable to \mathbf{CG} as a whole.

Intuitively, the next definition poses some global requirements for \mathbf{CG} to warrant its eligibility as a model-substitute.

Definition 4 (Coherence). Let $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$ be a completion graph for \mathcal{K} . \mathbf{CG} is called *coherent* iff

- for each $a \in \text{IN}_{\mathcal{K}}$ there is a unique element $\varepsilon_a \in \Delta$ such that $(v : a) \in \mathcal{S}(\varepsilon_a)$ for all variables v in $\mathcal{S}(\varepsilon_a)$,
- for each $\varepsilon, \varepsilon' \in \Delta$ and each variable v contained in $\mathcal{S}(\varepsilon)$, $\mathcal{S}(\varepsilon')$ contains some v' such that $\text{st}_{\varepsilon}(v) = \text{st}_{\varepsilon'}(v')$, and
- if $(v : \phi) \in \mathcal{S}(\varepsilon)$ and $\text{st}_{\varepsilon}(v) = \text{st}_{\varepsilon'}(v')$, then $(v' : \phi) \in \mathcal{S}(\varepsilon')$. \diamond

As usual in tableaux, inconsistencies emerge as clashes.

Definition 5 (Clash). A clash is a constraint of the form $(x : \perp)$. A completion graph \mathbf{CG} is said to contain a clash iff $\mathcal{S}(\varepsilon)$ does for some ε in \mathbf{CG} . Constraint systems or completion graphs not containing clashes are called *clash-free*. \diamond

The Algorithm

To decide whether a given $\mathbb{S}_{\mathcal{EL}}$ knowledge base \mathcal{K} in normal form is satisfiable, we form the initial completion graph \mathbf{CG}_I with $\mathcal{R} = \emptyset$ and Δ consisting of an element ε_{\top} with $\mathcal{L}(\varepsilon_{\top}) = \emptyset$ and $\mathcal{S}(\varepsilon_{\top}) = S_0^{\mathcal{K}}$, and for every $a \in \text{IN}_{\mathcal{K}}$ an element ε_a with $\mathcal{L}(\varepsilon_a) = \emptyset$ and $\mathcal{S}(\varepsilon_a) = S_0^{\mathcal{K}} \cup \{(x_s : a) \mid (x_s : \top) \in S_0^{\mathcal{K}}\}$.

After that, we repeatedly apply the local and global completion rules from Figure 3, where LL rules have the highest priority, followed by LC, GN, and GG rules, in that order. After each rule application, we check if \mathbf{CG} contains a clash and terminate with answer “unsatisfiable” should this be the case. If we arrive at a clash-free \mathbf{CG} with no more rules applicable, the algorithm terminates and returns “satisfiable”. The first table in Figure 2 illustrates the generated tableau for the $\mathbb{S}_{\mathcal{EL}}$ knowledge base $\mathcal{K} = \{\top \sqsubseteq \diamond_s C, C \sqsubseteq \exists R.D\}$.

Quasi-Models and Quasi-Satisfiability

This section sketches how special structures, called (dual) quasi-models can serve as proxies for proper $\mathbb{S}_{\mathcal{EL}}$ models.

Definition 6 (Run, Quasi-model). Let $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$ be a completion graph. A run r in \mathbf{CG} is a function mapping each element $\varepsilon \in \Delta$ to a variable of $\mathcal{S}(\varepsilon)$, such that

- C1 if $(r(\varepsilon) : s) \in \mathcal{S}(\varepsilon)$, then $(r(\varepsilon') : s) \in \mathcal{S}(\varepsilon')$ for all $\varepsilon, \varepsilon' \in \Delta$ and $s \in \text{ST}_{\mathcal{K}}$,
- C2 if $\langle \varepsilon, v, \varepsilon', v', R \rangle \in \mathcal{R}$ and $r(\varepsilon) = v$, then $r(\varepsilon') = v'$, and
- C3 if $(r(\varepsilon) : \exists R.C) \in \mathcal{S}(\varepsilon)$, there exists some $\varepsilon' \in \Delta$ with $\langle \varepsilon, r(\varepsilon), \varepsilon', r(\varepsilon'), R \rangle \in \mathcal{R}$ and $(r(\varepsilon') : C) \in \mathcal{S}(\varepsilon')$.

A quasi-model of \mathcal{K} is a tuple $\mathcal{Q} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R}, \Gamma \rangle$ where $\langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$ is a globally complete, coherent and clash-free completion graph for \mathcal{K} , and Γ a set of runs in $\langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$ such that for every $\varepsilon \in \Delta$ and variable v in $\mathcal{S}(\varepsilon)$, there is a run r in Γ such that $r(\varepsilon) = v$. \mathcal{K} is called *quasi-satisfiable* iff \mathcal{K} has a quasi-model. \diamond

In a nutshell, runs serve the purpose of lining up “compatible” variables, one from each individual constraint system, in a way that precisifications can be constructed (cf. Figure 2: in \mathcal{Q} , a “compatible” set of runs over \mathbf{CG} is displayed using dotted lines). With these notions in place, we can establish the desired result.

Theorem 2. A $\mathbb{S}_{\mathcal{EL}}$ knowledge base \mathcal{K} is satisfiable iff it is quasi-satisfiable.

Local labelling (LL) rule:

\mathbf{R}_{\preceq} If $\{x : s, x' : s \preceq s'\} \subseteq S$ but $(x : s') \notin S$, then set $S := S \cup \{x : s'\}$.

Local content (LC) rules:

\mathbf{R}_{\sqcap} If $\{x : C, x : D\} \subseteq S$, $(x : C \sqcap D) \notin S$ and $C \sqcap D \in \mathbf{C}_{\mathcal{K}}$, then set $S := S \cup \{x : C \sqcap D\}$.

\mathbf{R}_{\sqsubseteq} If $\{x : C, x : C \sqsubseteq D\} \subseteq S$ but $(x : D) \notin S$, then set $S := S \cup \{x : D\}$.

\mathbf{R}_{\square} If $\{x : \square_s \Phi, x' : s\} \subseteq S$ but $(x' : \Phi) \notin S$, then set $S := S \cup \{x' : \Phi\}$.

\mathbf{R}_g If $(x : \mathbf{G}) \in S$ but $(x' : \mathbf{G}) \notin S$, then set $S := S \cup \{x' : \mathbf{G}\}$.

\mathbf{R}_a If $\{x : a, x : C(a)\} \subseteq S$ but $(x : C) \notin S$, then set $S := S \cup \{x : C\}$.

\mathbf{R}_{\diamond} If $(x : \diamond_s C) \in S$ and $\{x' : s, x' : C\} \not\subseteq S$ for all x' in S , then create a fresh variable x' and set $S := S \cup \{x' : C, x' : s, x' : *, x' : \top\}$.

Global non-generating (GN) rules:

\mathbf{R}_{\downarrow} If $(x : C) \in \mathcal{S}(\varepsilon)$, $(\varepsilon', x', \varepsilon, x, R) \in \mathcal{R}$, and $\exists R.C \in \mathbf{C}_{\mathcal{K}}$, but $(x' : \exists R.C) \notin \mathcal{S}(\varepsilon')$, then set $\mathcal{S}(\varepsilon') := \mathcal{S}(\varepsilon') \cup \{x' : \exists R.C\}$.

\mathbf{R}_r If $\{x : a, x : R(a, b)\} \subseteq \mathcal{S}(\varepsilon)$ and $(x' : b) \in \mathcal{S}(\varepsilon')$, but $(\varepsilon, x, \varepsilon', x, R) \notin \mathcal{R}$, then set $\mathcal{S}(\varepsilon') := \mathcal{S}(\varepsilon') \cup \{x : \top\} \cup \{x : s \mid s \in \text{st}_{\varepsilon}(x)\}$ and $\mathcal{R} := \mathcal{R} \cup \{(\varepsilon, x, \varepsilon', x, R)\}$.

$\mathbf{R}_{r'}$ If $\{x : b, x : R(a, b)\} \subseteq \mathcal{S}(\varepsilon)$ and $(x' : a) \in \mathcal{S}(\varepsilon')$, but $(\varepsilon', x, \varepsilon, x, R) \notin \mathcal{R}$, then set $\mathcal{S}(\varepsilon') := \mathcal{S}(\varepsilon') \cup \{x : \top\} \cup \{x : s \mid s \in \text{st}_{\varepsilon}(x)\}$ and $\mathcal{R} := \mathcal{R} \cup \{(\varepsilon', x, \varepsilon, x, R)\}$.

\mathbf{R}_{\exists} If $(x : \exists R.C) \in \mathcal{S}(\varepsilon)$, $(C, \text{st}_{\varepsilon}(x), x') \in \mathcal{L}(\varepsilon')$ with $\varepsilon \neq \varepsilon'$ or $x = x'$, but $(\varepsilon, x, \varepsilon'', x'', R) \notin \mathcal{R}$ whenever $(C, \text{st}_{\varepsilon}(x), x'') \in \mathcal{L}(\varepsilon'')$ and $\varepsilon \neq \varepsilon''$ or $x = x''$, then set $\mathcal{R} := \mathcal{R} \cup \{(\varepsilon, x, \varepsilon', x', R)\}$.

Global generating (GG) rule:

\mathbf{R}_{\exists} If $(x : \exists R.C) \in \mathcal{S}(\varepsilon)$, but $(\varepsilon, x, \varepsilon'', x'', R) \notin \mathcal{R}$ whenever $(C, \text{st}_{\varepsilon}(x), x'') \in \mathcal{L}(\varepsilon'')$ and $\varepsilon \neq \varepsilon''$ or $x = x''$, then create ε' and a fresh variable x' , and then set $\mathcal{L}(\varepsilon') := \{(C, \text{st}_{\varepsilon}(x), x')\}$, $\mathcal{S}(\varepsilon') := \mathcal{S}_0^{\mathcal{K}} \cup \{x' : C, x' : \top\} \cup \{x' : s \mid s \in \text{st}_{\varepsilon}(x)\}$, $\mathcal{R} := \mathcal{R} \cup \{(\varepsilon, x, \varepsilon', x', R)\}$.

Figure 3: The tableau completion rules. \mathbf{G} can be of the form a , $(s \preceq s')$, or $\square_s \phi$. Φ may denote any element of $\text{SF}_{\mathcal{K}} \cup \mathbf{C}_{\mathcal{K}}$.

Proof. (sketch) We prove the correspondence by showing that every quasi-model gives rise to a model and vice versa.

(\Leftarrow) Given a quasi-model $\mathcal{Q} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R}, \Gamma \rangle$ of \mathcal{K} , we obtain a model $\mathcal{D} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ by letting $\Pi = \Gamma$, $\sigma(s) = \{r \mid (r(\varepsilon) : s) \in \mathcal{S}(\varepsilon)\}$, $C^{\gamma(r)} = \{\varepsilon \mid (r(\varepsilon) : C) \in \mathcal{S}(\varepsilon)\}$ for $C \in \text{BC}_{\mathcal{K}}$, $R^{\gamma(r)} = \{(\varepsilon, \varepsilon') \mid (\varepsilon, r(\varepsilon), \varepsilon', r(\varepsilon'), R) \in \mathcal{R}\}$ and $a^{\gamma(r)} = \varepsilon_a$ for all $a \in \text{IN}_{\mathcal{K}}$. (cf. in Figure 2, see \mathcal{D} and \mathcal{Q} .)

(\Rightarrow) Given a model $\mathcal{D} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ of \mathcal{K} , we obtain a quasi-model $\mathcal{Q} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R}, \Gamma \rangle$ as follows: Let $P := \Pi \cup \text{ST}_{\mathcal{K}}$ and for $p \in P$, let \bar{p} denote some arbitrary but fixed $\pi \in \sigma(p)$ if $p \in \text{ST}_{\mathcal{K}}$, and otherwise $\bar{p} = p$. Finally, let

$$\begin{aligned} \mathcal{S}(\varepsilon) &:= \{v_p : \phi \mid \mathcal{D}, \bar{p} \models \phi, \phi \in \text{SF}_{\mathcal{K}}\} \\ &\quad \cup \{v_p : C \mid \varepsilon \in C^{\gamma(\bar{p})}, p \in P\} \\ &\quad \cup \{v_p : a \mid \varepsilon = a^{\gamma(\bar{p})}, p \in P\} \cup \{v_p : s \mid \bar{p} \in \sigma(s)\} \\ \mathcal{L}(\varepsilon) &:= \{(C, \{s \mid \bar{p} \in \sigma(s)\}, v_p) \\ &\quad \mid \exists R.C \in \mathbf{C}_{\mathcal{K}}, (\varepsilon', \varepsilon) \in R^{\gamma(\bar{p})}, \varepsilon \in C^{\gamma(\bar{p})}\} \\ \mathcal{R} &:= \{(\varepsilon, v_p, \varepsilon', v_p, R) \mid (\varepsilon, \varepsilon') \in R^{\gamma(\bar{p})}\} \\ \Gamma &:= \{\{\varepsilon \mapsto v_p \mid \varepsilon \in \Delta\} \mid p \in P\}. \quad \square \end{aligned}$$

Polytime Termination and Correctness

Next, we give an overview of our argument why our algorithm runs in polynomial time with respect to $\|\mathcal{K}\|$, the size of its input \mathcal{K} . We observe that the number $|\Delta|$ of domain elements of any completion graph \mathbf{CG} constructed by our algorithm is bounded by $3\|\mathcal{K}\|^2$ (\dagger). We also find that the number of variables used in any single $\mathcal{S}(\varepsilon)$ is bounded by $2\|\mathcal{K}\|^2$ and the number of constraints in $\mathcal{S}(\varepsilon)$ by $2\|\mathcal{K}\|^3$ (\ddagger). Now, the number of applications of \mathbf{R}_{\exists} is bounded by the number of elements in each completion graph, i.e. at most $3\|\mathcal{K}\|^2$ in view of (\dagger). Since the rules \mathbf{R}_{\preceq} , \mathbf{R}_{\sqcap} , \mathbf{R}_{\sqsubseteq} , \mathbf{R}_{\square} , \mathbf{R}_{\diamond} , \mathbf{R}_g , \mathbf{R}_a , \mathbf{R}_r , $\mathbf{R}_{r'}$ and \mathbf{R}_{\downarrow} produce one or more new constraints in an element, the number of applications of such rules per element is bounded by $2\|\mathcal{K}\|^3$ due to (\ddagger). \mathbf{R}_{\exists} can add, for each ε with $(C, s, x) \in \mathcal{L}(\varepsilon)$, at most one quasi-role from every variable in every element, thus we have at most $6\|\mathcal{K}\|^4$ rule applications. The total number of rule applications is bounded by

the rule applications per element multiplied by the bound on elements, together with the bound on \mathbf{R}_{\exists} , which gives us

$$(6\|\mathcal{K}\|^4)(3\|\mathcal{K}\|^2) + (2\|\mathcal{K}\|^3)(3\|\mathcal{K}\|^2) + 3\|\mathcal{K}\|^2 \leq (27\|\mathcal{K}\|^6).$$

Theorem 3. *The completion algorithm terminates after at most $c\|\mathcal{K}\|^6$ steps, where c is a constant.*

As every single rule application can be clearly executed in polynomial time with respect to \mathcal{K} , we can conclude that our algorithm runs in polynomial time.

We are now ready to establish correctness of our decision algorithm, by showing its soundness and completeness. For both directions, Theorem 2 will come in handy. As usual, the soundness part of our argument is the easier one.

Theorem 4 (Soundness). *If there is a globally complete, coherent and clash-free completion graph \mathbf{CG} for a knowledge base \mathcal{K} , then \mathcal{K} is satisfiable.*

Proof. (sketch) Given $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$, let Γ consist of all runs on \mathbf{CG} . Then we can show that $\mathcal{Q} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R}, \Gamma \rangle$ constitutes a quasi-model for \mathcal{K} , so we can conclude by Theorem 2 that \mathcal{K} is satisfiable. \square

Proving completeness requires significantly more work. We make use of a notion that, intuitively, formalizes the idea that a completion graph \mathbf{CG} under development is “in sync” with a quasi-model \mathcal{Q} of the same knowledge base, where \mathcal{Q} can be conceived as a model-theoretic “upper bound” of \mathbf{CG} .

Definition 7 (Q-compatibility). *Let \mathcal{K} be a $\mathbb{S}_{\mathcal{E}\mathcal{L}}$ knowledge base and $\mathcal{Q} = \langle \Delta^q, \mathcal{S}^q, \mathcal{L}^q, \mathcal{R}^q, \Gamma^q \rangle$ be a quasimodel for \mathcal{K} . A completion graph $\mathbf{CG} = \langle \Delta, \mathcal{S}, \mathcal{L}, \mathcal{R} \rangle$ for \mathcal{K} is called Q-compatible iff there is a left-total relation $\mu \subseteq \Delta \times \Delta^q$ where*

- for all $g \in \Delta$ and $\varepsilon \in \Delta^q$, if both $\mathcal{L}(g) \subseteq \mathcal{L}^q(\varepsilon)$ and $\{a \mid (x : a) \in \mathcal{S}(g)\} \subseteq \{a \mid (v : a) \in \mathcal{S}^q(\varepsilon)\}$, then $(g, \varepsilon) \in \mu$,
- for each $(g, \varepsilon) \in \mu$ there is a surjective function $\mu_{g, \varepsilon}$ from the variables in $\mathcal{S}^q(\varepsilon)$ to the variables in $\mathcal{S}(g)$ such that
 - $(\mu_{g, \varepsilon}(v) : s) \in \mathcal{S}(g)$ implies $(v : s) \in \mathcal{S}^q(\varepsilon)$,
 - $(\mu_{g, \varepsilon}(v) : \Phi) \in \mathcal{S}(g)$ implies $(v : \Phi) \in \mathcal{S}^q(\varepsilon)$,
 - if $(g, x, g', x', R) \in \mathcal{R}$ then $(\varepsilon, y, \varepsilon', y', R) \in \mathcal{R}^q$ for some $(g, \varepsilon), (g', \varepsilon') \in \mu$ with $\mu_{g, \varepsilon}(y) = x$ and $\mu_{g', \varepsilon'}(y') = x'$. \diamond

With this definition, we can establish two important insights:

- The tableau algorithm’s initial completion graph \mathbf{CG}_I is \mathcal{Q} -compatible for any quasimodel \mathcal{Q} of \mathcal{K} .
- Applications of tableau rules preserve \mathcal{Q} -compatibility.

This entails that the completion graph maintained in the algorithm will be \mathcal{Q} -compatible at all times, thus also upon termination. We exploit this insight to show completeness.

Theorem 5 (Completeness). *If a $\mathbb{S}_{\mathcal{EL}}$ knowledge base \mathcal{K} is satisfiable, the tableau algorithm will construct a globally complete, coherent, and clash-free completion graph for \mathcal{K} .*

Proof. If \mathcal{K} is satisfiable then by Theorem 2, there is a quasimodel \mathcal{Q} for \mathcal{K} . According to Theorem 3, we can obtain a globally complete completion graph \mathbf{CG} after polynomially many applications of the tableau rules, which, as just discussed, is \mathcal{Q} -compatible. It must thus also be clash-free, because otherwise there were an element g and variable x , with $(x : \perp) \in \mathcal{S}(g)$, and thus there is $(g, \varepsilon) \in \mu$ and $\mu_{g, \varepsilon}$ such that $(\mu_{g, \varepsilon}(x) : \perp) \in \mathcal{S}^q(\varepsilon)$, which is a contradiction because \mathcal{Q} is a quasi-model. It is not hard to show that \mathbf{CG} is also coherent, whence we can conclude that \mathbf{CG} is a globally complete, coherent, and clash-free completion graph for \mathcal{K} . \square

Together with the well-known PTIME-hardness of the satisfiability problem in (standpoint-free) \mathcal{EL} , we have therefore established PTIME-completeness of $\mathbb{S}_{\mathcal{EL}}$ and exhibited a worst-case optimal algorithm for it.

4 Intractable Extensions

While the shown tractability of reasoning in $\mathbb{S}_{\mathcal{EL}}$ is good news, one might ask if one could include more modelling features or relax certain side conditions and still preserve tractability. This section shows that tractability can be easily lost (at least under standard complexity-theoretic assumptions).

Empty standpoints

While it may make sense on a philosophical level, one might wonder whether the constraint that $\sigma(s)$ needs to be nonempty for every $s \in \text{ST}_{\mathcal{K}}$ has an impact on tractability. In fact, dropping this constraint, obtaining a logic $\mathbb{S}_{\mathcal{EL}}^0$ with the same syntax but modified semantics, would increase expressivity (standpoint non-emptiness could still be enforced in $\mathbb{S}_{\mathcal{EL}}^0$ by asserting $\top \sqsubseteq \diamond_s \top$ for every $s \in \text{ST}_{\mathcal{K}}$). However, satisfiability in $\mathbb{S}_{\mathcal{EL}}^0$ turns out to be NP-hard, even when disallowing usage of concept and role names entirely. The key insight that both $\diamond_s \top$ and its negation $\square_s \perp$ can be expressed as $\mathbb{S}_{\mathcal{EL}}^0$ concepts gives rise to the following reduction from 3SAT: Assume an instance $\phi = \bigvee \mathbf{C}_1 \wedge \dots \wedge \bigvee \mathbf{C}_n$ of 3SAT containing n clauses (i.e., disjunctions of literals) $\bigvee \mathbf{C}_j$ over the propositional variables $P = \{p_1, \dots, p_k\}$. We note that ϕ is equivalent to $(\bigwedge \overline{\mathbf{C}}_1 \rightarrow \text{false}) \wedge \dots \wedge (\bigwedge \overline{\mathbf{C}}_n \rightarrow \text{false})$, where $\overline{\mathbf{C}}_j$ is obtained from \mathbf{C}_j by replacing every literal by its negated version. Let now $\{s_1, \dots, s_k\}$ be a set of standpoint names and, for any literal ℓ over P , define

$$L_\ell = \begin{cases} \diamond_{s_i} \top & \text{if } \ell = p_i, \\ \square_{s_i} \perp & \text{if } \ell = \neg p_i. \end{cases}$$

Then, ϕ is satisfiable iff the following $\mathbb{S}_{\mathcal{EL}}^0$ knowledge base is: $\mathcal{K}_\phi = \{L_\ell \sqcap L_{\ell'} \sqcap L_{\ell''} \sqsubseteq \perp \mid \{\ell, \ell', \ell''\} = \overline{\mathbf{C}}_j, 1 \leq j \leq n\}$.

Rigid roles

$\mathbb{S}_{\mathcal{EL}}$ allows knowledge engineers to globally enforce rigidity of specific concepts through axioms of the shape $A \sqsubseteq \square_* A$. (This is in contrast to e.g. $\mathbf{K}_n \mathcal{ALC}$, where rigidity of concepts can only be expressed *relative* to a given formula.) In a similar manner, rigidity of roles (i.e., the interpretation of certain distinguished roles being the same throughout all precisifications) would represent a desirable modelling feature. Other modal extensions of DLs have easily been shown to even become undecidable when this feature is permitted, but as $\mathbb{S}_{\mathcal{EL}}$ uses a much simplified semantics on the modal dimension, these results do not carry over to $\mathbb{S}_{\mathcal{EL}}$. Yet, we will show that just the presence of *one* distinguished rigid role \dot{R} causes $\mathbb{S}_{\mathcal{EL}}$ to become intractable as satisfiability turns CONP-hard. To demonstrate this, we reduce 3SAT to KB unsatisfiability. As above, assume an instance $\phi = \bigvee \mathbf{C}_1 \wedge \dots \wedge \bigvee \mathbf{C}_n$ of 3SAT over propositional variables $P = \{p_1, \dots, p_k\}$. Then ϕ is satisfiable iff the following $\mathbb{S}_{\mathcal{EL}}$ TBox is unsatisfiable (with all axioms instantiated for $1 \leq i \leq k$):

$$\begin{aligned} \top &\sqsubseteq \exists T.L_0 \\ L_{i-1} &\sqsubseteq \exists \dot{R}.\square_*(L_i \sqcap T_{p_i}) \\ L_{i-1} &\sqsubseteq \exists \dot{R}.\square_*(L_i \sqcap T_{\neg p_i}) \\ L_k &\sqsubseteq \diamond_* S \\ \exists \dot{R}.(T_{p_i} \sqcap S) &\sqsubseteq (T_{p_i} \sqcap S) \\ \exists \dot{R}.(T_{\neg p_i} \sqcap S) &\sqsubseteq (T_{\neg p_i} \sqcap S) \\ L_0 \sqcap T_\ell &\sqsubseteq T_{\mathbf{C}_j} \text{ for all } \ell \in \mathbf{C}_j \\ T_{\mathbf{C}_1} \sqcap \dots \sqcap T_{\mathbf{C}_n} &\sqsubseteq \perp \end{aligned}$$

The intuition behind this is to construct a “decision tree” where \dot{R} acts as child relation and which (thanks to \dot{R} ’s rigidity) is synchronized across all precisifications. The leafs of this tree correspond to all possible truth assignments for p_1, \dots, p_k . Then we make sure that for every leaf, there exists a precisification, where this leaf is “selected” (indicated by the non-rigid concept S). This “selection marker” S is propagated from children to parents, taking along information on the truth assignments. Finally, when all the information on the path leading to the selected node has been accumulated in the root node, it is checked if this information corresponds to an assignment satisfying ϕ . If so, a “global inconsistency” is triggered. Figure 4 provides a small example.




Figure 4: Standpoint structure, witnessing the unsatisfiability of the propositional formula $\phi = (\neg p_1 \vee p_2) \wedge (p_1) \wedge (\neg p_2)$. All arrows indicate \dot{R} ; role T is omitted for better readability. The concept names displayed in the top row hold throughout all precisifications.

Nominal Concepts

Nominal concepts are a modelling feature widely used in ontology languages. For an individual o , the nominal concept $\{o\}$ refers to the singleton set $\{o^{\mathcal{T}}\}$. Let $\mathcal{EL}\mathcal{O}$ denote \mathcal{EL} extended by nominal concepts. Several formalisms subsuming $\mathcal{EL}\mathcal{O}$, including OWL 2 EL, are known to allow for tractable reasoning [Baader *et al.*, 2005; Krötzsch, 2010]. However, in the presence of standpoints, nominals prove to be detrimental for the reasoning complexity: satisfiability of $\mathbb{S}_{\mathcal{EL}\mathcal{O}}$ TBoxes using just one nominal concept $\{o\}$ turns out to be EXPTIME-hard and thus definitely harder than for $\mathbb{S}_{\mathcal{EL}}$. This can be shown by a PTIME reduction of satisfiability for Horn- \mathcal{ALC} TBoxes (which is known to be EXPTIME-complete [Krötzsch *et al.*, 2013]) to satisfiability of $\mathbb{S}_{\mathcal{EL}\mathcal{O}}$ TBoxes with just one standpoint (the global one) and one nominal concept $\{o\}$. To this end, recall that any Horn- \mathcal{ALC} TBox can be normalised in PTIME to consist of only axioms of the following shapes:

$$A \sqsubseteq B \quad A \sqcap B \sqsubseteq C \quad \exists R.A \sqsubseteq B \quad A \sqsubseteq \exists R.B \quad A \sqsubseteq \forall R.B$$

where A, B, C can be concept names, \top , or \perp . From a normalised Horn- \mathcal{ALC} TBox \mathcal{T} , we obtain the target $\mathbb{S}_{\mathcal{EL}\mathcal{O}}$ TBox \mathcal{T}' by (i) declaring every original concept name as rigid via the axiom $A \sqsubseteq \square_* A$ as well as (ii) replacing every axiom of the shape $A \sqsubseteq \exists R.B$ by the axiom

$$A \sqsubseteq \diamond_* ((\exists \text{Src}.\{o\}) \sqcap (\exists R.(B \sqcap \exists \text{Tgt}.\{o\})))$$

(introducing two fresh role names Src and Tgt), and replacing every axiom of the shape $A \sqsubseteq \forall R.B$ by the two axioms

$$A \sqcap \exists R.\top \sqsubseteq (\exists \text{Src}.\{o\} \sqcap \tilde{B}) \quad \text{and} \quad \exists \text{Tgt}.\tilde{B} \sqsubseteq B,$$

introducing a (non-rigid) copy \tilde{A} for every original concept name A . With this polytime translation, satisfiability of the Horn- \mathcal{ALC} TBox \mathcal{T} and the $\mathbb{S}_{\mathcal{EL}\mathcal{O}}$ TBox \mathcal{T}' coincide.

The intuition behind this is to emulate the DL-model of \mathcal{T} by a standpoint structure where every role connection gets assigned its own precisification, wherein the nominal acts as “witness” for this connection and is linked to source and target element by dedicated roles. Then the “forward transfer” of information by axioms of the shape $A \sqsubseteq \forall R.B$ can be realized using the nominal element as “proxy” (cf. Figure 5).

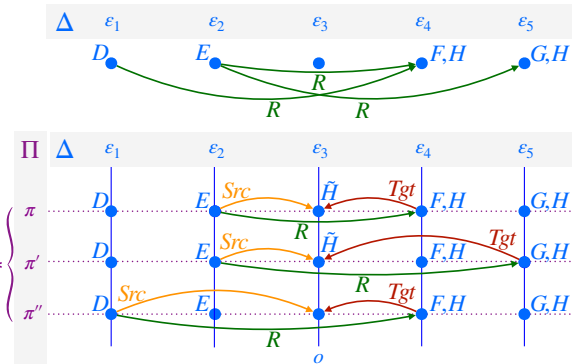


Figure 5: Top: DL interpretation satisfying the Horn- \mathcal{ALC} TBox $\mathcal{T} = \{D \sqsubseteq \exists R.\top, E \sqsubseteq \exists R.F, E \sqsubseteq \exists R.G, E \sqsubseteq \forall R.H\}$. Bottom: Corresponding standpoint structure satisfying \mathcal{T}' .

5 Conclusion and Future Work

In this paper, we introduced Standpoint \mathcal{EL} , a new lightweight member of the emerging family of standpoint logics. We described the new modelling and reasoning capabilities it brings to large-scale ontology management and established a PTIME (and thus worst-case optimal) tableau-based decision procedure for standard reasoning tasks. We also demonstrated that certain extensions of $\mathbb{S}_{\mathcal{EL}}$, which would be desirable from an expressivity point of view, inevitably come with a loss of tractability (sometimes under the assumption $P \neq NP$).

Yet, several modelling features can be accommodated into $\mathbb{S}_{\mathcal{EL}}$ without endangering tractability. For instance, from a practical perspective, it appears very desirable and advantageous modelling-wise, if not just single axioms, but whole axiom sets (up to whole knowledge bases) could be preceded by standpoint modalities. By definition, an axiom of the type $\square_s \mathcal{K}$ can be equivalently rewritten into the axiom set $\{\square_s \phi \mid \phi \in \mathcal{K}\}$. While something alike is not immediately possible for axioms of the type $\diamond_s \mathcal{K}$, our normalization rule for diamond-preceded axioms can be lifted and thus $\diamond_s \mathcal{K}$ can be rewritten to $\square_{s'} \mathcal{K}$ (and further to $\{\square_{s'} \phi \mid \phi \in \mathcal{K}\}$) upon introducing a fresh standpoint name s' and asserting $s' \preceq s$. Thus standpoint-modality-annotated knowledge bases come essentially for free in $\mathbb{S}_{\mathcal{EL}}$. In fact, we already made tacit use of this modelling feature in Axiom 9 and Axiom 10 of our initial example.

Moreover, we are confident that, as opposed to nominal concepts, other modelling features of OWL 2 EL can be added to $\mathbb{S}_{\mathcal{EL}}$ without harming tractability. These include complex role inclusions (also called role-chain axioms) such as $\text{FindingSite} \circ \text{PartOf} \sqsubseteq \text{FindingSite}$, and the self-concept as in $\text{ApoptoticCell} \sqsubseteq \exists \text{Destroys}.\text{Self}$.

Beyond exploring the tractability boundaries, potential next research endeavours include to investigate diverse feasible strategies for developing a $\mathbb{S}_{\mathcal{EL}}$ reasoner. Options worth pursuing toward this goal include

- to implement our tableau algorithm from scratch or by modifying existing open-source tableaux systems,
- to design a deduction calculus over normalised axioms that can be translated into a datalog program, akin to the approach of Krötzsch [2010], then utilizing a state-of-the-art datalog engine like VLog [Urbani *et al.*, 2016], or
- to find a reduction to reasoning in standpoint-free (PTIME extensions of) \mathcal{EL} that is supported by existing reasoners (such as ELK [Kazakov *et al.*, 2014]).

With reasoners in place, appropriate experiments can be conducted to assess practical feasibility and scalability.

In addition to the \mathcal{EL} family, further popular and computationally lightweight formalisms exist, such as the tractable profiles OWL 2 RL and OWL 2 QL [Motik *et al.*, 2009a]. It would be interesting to investigate options to extend these by standpoint reasoning without sacrificing tractability. More generally, we intend to research the effect of adding standpoints to KR languages – light- or heavyweight – in terms of computational properties and expressivity as well as avenues for implementing efficient reasoners for them. Beyond the large and versatile family of description logics, a worthwhile target for these efforts would be existential rule languages.

Acknowledgements This work was supported by funding from BMBF within project *KIMEDS* (grant no. GW0552B), by the European Research Council through the ERC Consolidator Grant *DeciGUT* (771779), and by the Federal Ministry of Education and Research of Germany and the Sächsische Staatsministerium für Wissenschaft Kultur und Tourismus through the Center of Excellence for AI-research *ScaDS.AI*.

References

- [Baader and Ohlbach, 1995] Franz Baader and Hans Jürgen Ohlbach. A multi-dimensional terminological knowledge representation language. *Journal of Applied Non-Classical Logics*, 5(2):153–197, 1995.
- [Baader et al., 2005] Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the EL envelope. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, 2005*, pages 364–369. Professional Book Center, 2005.
- [Baader et al., 2017] Franz Baader, Ian Horrocks, Carsten Lutz, and Uli Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.
- [Borgida and Serafini, 2003] Alex Borgida and Luciano Serafini. Distributed Description Logics: Assimilating information from peer sources. *Journal on Data Semantics*, 2800:153–184, 2003.
- [Bouquet et al., 2003] Paolo Bouquet, Fausto Giunchiglia, Frank Van Harmelen, Luciano Serafini, and Heiner Stuckenschmidt. C-OWL: Contextualizing ontologies. *Lecture Notes in Computer Science*, 2870:164–179, 10 2003.
- [Bozzato et al., 2018] Loris Bozzato, Thomas Eiter, and Luciano Serafini. Enhancing context knowledge repositories with justifiable exceptions. *Artificial Intelligence*, 257:72–126, 2018.
- [Carral et al., 2014] David Carral, Adila Krisnadhi, Sebastian Rudolph, and Pascal Hitzler. All but not nothing: Left-hand side universals for tractable OWL profiles. In C. Maria Keet and Valentina A. M. Tamma, editors, *Proceedings of the 11th International Workshop on OWL: Experiences and Directions (OWLED 2014) co-located with 13th International Semantic Web Conference on (ISWC 2014)*, volume 1265, pages 97–108. CEUR Workshop Proceedings, 2014.
- [Donnelly, 2006] Kevin Donnelly. SNOMED-CT: The advanced terminology and coding system for eHealth. *Studies in health technology and informatics*, 121:279, 2006.
- [Fine, 1975] Kit Fine. Vagueness, truth and logic. *Synthese*, 30(3-4):265–300, 1975.
- [Gómez Álvarez and Rudolph, 2021] Lucía Gómez Álvarez and Sebastian Rudolph. Standpoint logic: Multi-perspective knowledge representation. In Fabian Neuhaus and Boyan Brodaric, editors, *Proceedings of the 12th International Conference on Formal Ontology in Information Systems*, volume 344 of *FAIA*, pages 3–17. IOS Press, 2021.
- [Gómez Álvarez et al., 2022] Lucía Gómez Álvarez, Sebastian Rudolph, and Hannes Strass. How to Agree to Disagree: Managing Ontological Perspectives using Standpoint Logic. In *Proceedings of the 21st International Semantic Web Conference*, pages 125–141. Springer, 2022.
- [Hemam and Boufaïda, 2011] Mounir Hemam and Zizette Boufaïda. MVP-OWL: A multi-viewpoints ontology language for the Semantic Web. *International Journal of Reasoning-based Intelligent Systems*, 3(3-4):147–155, 2011.
- [Hemam, 2018] Mounir Hemam. An extension of the ontology web language with multi-viewpoints and probabilistic reasoning. *International Journal of Advanced Intelligence Paradigms*, 10(3):247–265, 2018.
- [Kazakov et al., 2014] Yevgeny Kazakov, Markus Krötzsch, and Frantisek Simancik. The incredible ELK - from polynomial procedures to efficient reasoning with \mathcal{EL} ontologies. *J. Autom. Reason.*, 53(1):1–61, 2014.
- [Klarman and Gutiérrez-Basulto, 2013] Szymon Klarman and Víctor Gutiérrez-Basulto. Description logics of context. *Journal of Logic and Computation*, 26(3):817–854, 2013.
- [Krötzsch et al., 2013] Markus Krötzsch, Sebastian Rudolph, and Pascal Hitzler. Complexities of Horn description logics. *ACM Trans. Comput. Log.*, 14(1):2:1–2:36, 2013.
- [Krötzsch, 2010] Markus Krötzsch. Efficient inferencing for OWL EL. In Tomi Janhunen and Ilkka Niemelä, editors, *Proceedings of the 12th European Conference on Logics in Artificial Intelligence JELIA 2010.*, volume 6341 of *Lecture Notes in Computer Science*, pages 234–246. Springer, 2010.
- [Lutz et al., 2002] Carsten Lutz, Holger Sturm, Frank Wolter, and Michael Zakharyashev. A tableau decision algorithm for modalized ALC with constant domains. *Studia Logica*, 72(2):199–232, 2002.
- [McCarthy and Buvac, 1998] John McCarthy and Sasa Buvac. Formalizing context (expanded notes). *CSLI Lecture Notes*, 81:13–50, 1998.
- [Mosurović, 1999] Milenko Mosurović. On the complexity of description logics with modal operators. 1999.
- [Motik et al., 2009a] Boris Motik, Bernardo Cuenca Grau, Ian Horrocks, Zhe Wu, Achille Fokoue, and Carsten Lutz, editors. *OWL 2 Web Ontology Language: Profiles*. W3C Recommendation, 27 October 2009. Available at <http://www.w3.org/TR/owl2-profiles/>.
- [Motik et al., 2009b] Boris Motik, Peter F. Patel-Schneider, and Bernardo Cuenca Grau, editors. *OWL 2 Web Ontology Language: Direct Semantics*. W3C Recommendation, 27 October 2009. Available at <http://www.w3.org/TR/owl2-direct-semantics/>.
- [Osman et al., 2021] Inès Osman, Sadok Ben Yahia, and Gayo Diallo. Ontology Integration: Approaches and Challenging Issues. *Information Fusion*, 71:38–63, 2021.

- [OWL Working Group, 2009] W3C OWL Working Group. *OWL 2 Web Ontology Language: Document Overview*. W3C Recommendation, 27 October 2009. Available at <http://www.w3.org/TR/owl2-overview/>.
- [Rudolph, 2011] Sebastian Rudolph. Foundations of description logics. In Axel Polleres, Claudia d’Amato, Marcelo Arenas, Siegfried Handschuh, Paula Kroner, Sascha Ossowski, and Peter F. Patel-Schneider, editors, *Lecture Notes of the 7th International Reasoning Web Summer School (RW’11)*, volume 6848 of *LNCS*, pages 76–136. Springer, 2011.
- [Schulz *et al.*, 2017] Stefan Schulz, Catalina Martínez-Costa, and Jose Antonio Miñarro-Giménez. Lexical ambiguity in SNOMED CT. In *JOWO*, 2017.
- [Serafini and Homola, 2012] Luciano Serafini and Martin Homola. Contextualized knowledge repositories for the semantic web. *Journal of Web Semantics*, 12-13:64–87, 2012. Reasoning with context in the Semantic Web.
- [Urbani *et al.*, 2016] Jacopo Urbani, Cerial J. H. Jacobs, and Markus Krötzsch. Column-oriented datalog materialization for large knowledge graphs. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 258–264. AAAI Press, 2016.
- [Wolter and Zakharyashev, 1998] Frank Wolter and Michael Zakharyashev. Satisfiability problem in description logics with modal operators. In Anthony G. Cohn, Lenhart K. Schubert, and Stuart C. Shapiro, editors, *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning*, pages 512–523. Morgan Kaufmann, 1998.
- [Wolter and Zakharyashev, 1999] Frank Wolter and Michael Zakharyashev. Multi-dimensional description logics. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI’99*, page 104–109. Morgan Kaufmann Publishers Inc., 1999.