Research Memorandum No. 1230 May 21st, 2026

BoxLitE: A faithful knowledge base embedding based on convex optimization

Bruno F. Lourenço (The Institute of Statistical Mathematics, Japan)
Hesham Morgan (TU Wien, Austria)
Ana Ozaki (University of Oslo, Norway)
Aleksandar Pavlović (University of Applied Sciences Campus Vienna, Austria)
Emanuel Sallinger (TU Wien, Austria)

The Institute of Statistical Mathematics

10-3 Midori-cho, Tachikawa,Tokyo, 190-8562, Japan

BoxLitE: A Faithful Knowledge Base Embedding Based on Convex Optimization

Bruno F. Lourenço[1] , **Hesham Morgan**[2] , **Ana Ozaki**[3] , **Aleksandar Pavlović**[4] , **Emanuel Sallinger**[2]
[1]The Institute of Statistical Mathematics, Japan
[2]TU Wien, Austria
[3]University of Oslo, Norway
[4]University of Applied Sciences Campus Vienna, Austria
bruno@ism.ac.jp, hesham.morgan@tuwien.ac.at, anaoz@uio.no, aleksandar.pavlovic@hcw.ac.at, sallinger@dbai.tuwien.ac.at

Abstract

Knowledge base (KB) embeddings aim at combining the capability of classical knowledge graph embeddings to generalize the information present in facts, the ABox, with conceptual knowledge represented in an ontology language, the TBox. Several authors have recently explored the idea of mapping concepts to *convex regions* in a vector space. This is useful to represent hierarchies, typically present in TBoxes, since more general concepts can be mapped to larger regions, containing those regions associated with more specific concepts. However, the power of convexity is rarely leveraged during the actual learning tasks. Here, we introduce BoxLitE, a KB embedding model for DL-Lite$^{\mathcal{H}}$ that allows for convex optimization. We show that for any satisfiable DL-Lite$^{\mathcal{H}}$ KB, there is a BoxLitE embedding that is a weakly faithful model. As a proof of concept, we show how to formulate the KB embedding task as a convex optimization problem and how to obtain embeddings with such desirable faithfulness property.

1 Introduction

Knowledge base (KB) embeddings combine the capability of knowledge graph embeddings to perform inductive reasoning for link prediction with deductive reasoning, using logic expressions present in an ontology (Bourgaux et al., 2024). Several authors have recently explored the idea of mapping concepts in a KB to regions in a vector space (Gutiérrez-Basulto and Schockaert, 2018; Pavlović and Sallinger, 2023a; Pavlović, Sallinger, and Schockaert, 2025; Morgan et al., 2026). Region-based embeddings are important for KBs as they can naturally represent hierarchies: more general concepts can be mapped to larger regions, containing those regions associated with more specific concepts.

Although concepts are usually mapped to *convex regions* in KBEs, e.g., balls, boxes, and cones (Kulmanov et al., 2019; Abboud et al., 2020; Xiong et al., 2022; Pavlović and Sallinger, 2023b; Lütfü Özçep, Leemhuis, and Wolter, 2020), the power of convexity is rarely leveraged during the actual learning tasks. While being convex is not synonymous with "easy", convexity brings a number of theoretical and practical benefits, in particular: all local minima must be global. Under convexity, there are a number of efficient algorithms for many classes of convex problems such as linear programs and second-order cone programs (SOCPs) (Nesterov and Nemirovskii, 1994; Ben-Tal and Nemirovski, 2001).

For KBs expressed in description logic, most of the literature work on region-based embeddings focuses on the $\mathcal{EL}$ ontology language (Yang, Chen, and Sattler, 2025; Jackermeier, Chen, and Horrocks, 2024; Lacerda, Ozaki, and Guimarães, 2024a; Xiong et al., 2022; Kulmanov et al., 2019; Mondal, Bhatia, and Mutharaju, 2021; Peng et al., 2022; Lacerda, Ozaki, and Guimarães, 2024b), while some consider $\mathcal{ALC}$ (Lütfü Özçep, Leemhuis, and Wolter, 2020; Leemhuis, Özçep, and Wolter, 2022). However, in practice, ontologies present in large-scale KBs tend to use features in the DL-Lite$^{\mathcal{H}}$ ontology language (Artale et al., 2009), due to its simple but versatile expressivity, featuring low computational complexity (Artale et al., 2009). Despite its broad use, only a few works investigate logics in the DL-Lite family in a KB embedding context (Li et al., 2022; Imenes, Guimarães, and Ozaki, 2023; Bourgaux, Ozaki, and Pan, 2021) and none provide a region-based embedding implementation.

In this paper, we introduce BoxLitE, a KB embedding model for DL-Lite$^{\mathcal{H}}$ KBs that allows for convex optimization. In particular, we study the notion of *weakly faithful models* (Lütfü Özçep, Leemhuis, and Wolter, 2020; Bourgaux et al., 2024) in the optimization task. This notion states that axioms that hold in the embedding are *consistent* with the KB and *entailments of the KB are satisfied* in the embedding. In detail,

- we select DL-Lite$^{\mathcal{H}}$ an **ontology language** that allows to exploit the **advantages of convex optimization**;
- we **design, implement, and evaluate** BoxLitE, a KB embedding approach that allows for convex optimization;
- we prove that every satisfiable DL-Lite$^{\mathcal{H}}$ has a BoxLitE embedding that is a **weakly faithful model**;
- we introduce a novel and efficient form of **negative sampling based on existential concepts**;
- in contrast to commonly used unconstrained optimization approaches, our BoxLitE approach can enforce the **TBox axioms using convex constraints** while leaving the ABox terms in the objective function.

Main Contribution. Our main contribution lies on theoretical grounds, establishing the existence of a faithful model for DL-Lite$^{\mathcal{H}}$ and the proposal of a novel approach that computes embeddings for DL-Lite$^{\mathcal{H}}$ KBs. BoxLitE's implementation and the empirical results serve as a proof of concept that our theoretical results translate into practical settings. One of the motivations for this work is to try to solve link prediction using theoretically sounds techniques from a convex optimization perspective. We would like to check how much performance can we get from a purely convex approach. In this way, our work stands in contrast to previous approaches that relied on nonconvexity, yielding more complicated optimization problems.

There are three primary sources of nonconvexity in contemporary KB embedding methods (see more details in Section 7): the way *negative sampling* is used by most works (Bordes et al., 2013; Sun et al., 2019; Yang et al., 2015; Trouillon et al., 2016; Kazemi and Poole, 2018; Balazevic, Allen, and Hospedales, 2019; Xiong et al., 2022; Abboud et al., 2020; Pavlović and Sallinger, 2024b); the *design of loss terms* involving operations that do not preserve convexity; and the choice of the *ontology language*, in particular, languages that allow conjunction on the left-hand side of concept inclusions, such as $\mathcal{EL}$ and $\mathcal{ALC}$, are likely to lead to nonconvexity. This motivated us to consider DL-Lite$^{\mathcal{H}}$ (Artale et al., 2009), which is a simple but practically useful and well-known ontology language without conjunction on the left-hand side.

On the algorithmic side, ADAM (Kingma and Ba, 2015), the method of choice of several works, is often used in theoretically unsound ways. For example, ADAM requires the objective function to be differentiable; however, it is not uncommon to see this requirement being ignored. For instance, in (Xiong et al., 2022, Sections 4.3, 4.5) and (Abboud et al., 2020, Section 4), the authors describe nondifferentiable objective functions that are optimized via ADAM[1]. The rationale for that is not explained in the papers, but a reasonable guess seems to be that these functions are typically differentiable almost everywhere (in a measure theory sense), so there may be an underlying belief that iterates are unlikely to reach a point of nondifferentiability and the algorithm may still work in practice. Besides, widely used tools such as PyTorch may attempt to differentiate nondifferentiable functions for the user by, e.g., selecting subgradients/supgradients when available (PyTorch, 2026). Unfortunately, there are well-known examples in the optimization literature showing that when a gradient method is applied directly to a nondifferentiable function, it may fail to find an optimal solution, even if the function is differentiable at the points generated by the method, e.g., see (Beck, 2017, Section 8.1.2). In the optimization community, there are works that explore the convergence properties of simple SGD methods when applied to nondifferentiable functions (Bolte and Pauwels, 2020; Davis et al., 2019), but, as far as we know, similar results have not been proved for ADAM.

When performance is good, one may be tempted to overlook these issues, but when it is not, it may be hard to know what is to blame. Is it the method, the parameter choice, the optimization algorithm, or a combination of those? Our proposed approach *does* include nondifferentiable terms as well, but, in contrast, our method of choice for solving the underlying optimization problem is suitable for handling nondifferentiability (Section 7). In this sense, our approach is conceptually sound from an optimization point of view.

Organization. Our paper is organized as follows. Section 2 provides basic definitions. Section 3 defines the semantics of our BoxLitE approach. Section 4 studies BoxLitE's faithfulness properties. Section 5 formulates BoxLitE's convex optimization problem and shows faithfulness properties that are ensured by the problem formulation. Section 6 discusses BoxLitE's proof of concept implementation and experiments. Section 8 provides a discussion on optimization in KB embeddings and concludes our paper. Omitted proofs and additional details are given in the appendix.

2 Basic Definitions: DL-Lite Ontologies

Let $\mathsf{N_C}$, $\mathsf{N_R}$, and $\mathsf{N_I}$ be *finite*, non-empty, mutually disjoint sets of *concept*, *role*, and *individual* names, respectively. We denote by $\mathsf{N_R^-}$ the set $\mathsf{N_R} \cup \{R^- \mid R \in \mathsf{N_R}\}$, by $\mathsf{N_C^\exists}$ the set $\mathsf{N_C} \cup \{\exists R \mid R \in \mathsf{N_R^-}\}$, and by $\mathsf{N_{C\neg}^\exists}$ the set $\mathsf{N_C^\exists} \cup \{\neg E \mid E \in \mathsf{N_C^\exists}\}$. DL-Lite$^{\mathcal{H}}$ role and concept inclusions are of the form $S \sqsubseteq T$ and $B \sqsubseteq C$, resp., where $S, T \in \mathsf{N_R^-}$ are roles[2] and B, C are concepts built as follows: $S ::= R \mid R^-,\ B ::= A \mid \exists S,\ C ::= B \mid \neg B$, with $R \in \mathsf{N_R}$ and $A \in \mathsf{N_C}$. A DL-Lite$^{\mathcal{H}}$ *TBox* (we may also use the less technical term *ontology*) is a (finite) set of DL-Lite$^{\mathcal{H}}$ concept and role inclusions. *Assertions* are of the form $D(a)$ (called *concept assertions*) or $R(a, b)$ (called *role assertions*), where $D \in \mathsf{N_C^\exists}$, $R \in \mathsf{N_R}$, and $a, b \in \mathsf{N_I}$. A DL-Lite$^{\mathcal{H}}$ *knowledge base* (KB) is a pair $(\mathcal{T}, \mathcal{A})$ where $\mathcal{T}$ is a DL-Lite$^{\mathcal{H}}$ TBox and $\mathcal{A}$ is a (finite) set of assertions, called *ABox*. Following the KBE literature (e.g. (Xiong et al., 2022)) and to simplify our presentation, we assume that DL-Lite$^{\mathcal{H}}$ TBoxes are in *named form*, meaning that they contain only concept inclusions where one of the concepts is a concept name. The semantics is given as usual by interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ (see appendix). We call a DL-Lite$^{\mathcal{H}}$ *axiom* an expression that is a role inclusion (RI), a concept inclusion (CI), or an assertion. We write $\mathcal{I} \models \alpha$ if $\mathcal{I}$ satisfies an axiom α. An interpretation $\mathcal{I}$ satisfies a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, written $\mathcal{I} \models \mathcal{K}$, if it satisfies the axioms in $\mathcal{T}$ and $\mathcal{A}$. We say that $\mathcal{K}$ is *satisfiable* if such interpretation $\mathcal{I}$ exists. Also, $\mathcal{K}$ *entails* an axiom α, written $\mathcal{K} \models \alpha$, iff, for all interpretations $\mathcal{I}$, if $\mathcal{I} \models \mathcal{K}$ then $\mathcal{I} \models \alpha$. We say that an axiom α is *consistent* with a KB $\mathcal{K}$ if there is an interpretation $\mathcal{I}$ such that $\mathcal{I} \models \mathcal{K} \cup \{\alpha\}$.

Canonical Model. Our construction is based on previous definitions of canonical models for DL-Lite$^{\mathcal{H}}$ (e.g., (Kontchakov et al., 2010)). Though here, we ensure that *inclusions* hold in the model *only if* they are entailed by the ontology. Before we

[1] In (Xiong et al., 2022), for example, loss terms for concepts assertions are expressed using the 2-norm and the regularization term considered therein use a maximum of functions. Both correspond to terms that are not differentiable in general, e.g., the 2-norm function is not differentiable at the origin, which can be checked directly by the definition of Fréchet differentiability.

[2] A *role* is a role name or the inverse of a role name.

provide the definition of the canonical model, we introduce the following notions. Assume $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is a satisfiable DL-Lite$^{\mathcal{H}}$ KB. We say that a concept C is *satisfiable w.r.t.* $\mathcal{K}$ if there is an interpretation $\mathcal{I}$ such that $\mathcal{I} \models \mathcal{K}$ and $C^{\mathcal{I}} \neq \emptyset$. We write concepts of the form $B \sqcap C$ (with the *conjunction* operator) only to falsify concept inclusions of the form $B \sqsubseteq \neg C$. In an interpretation $\mathcal{I}$, the meaning of $(B \sqcap C)^{\mathcal{I}}$ is $B^{\mathcal{I}} \cap C^{\mathcal{I}}$. Let $\mathsf{N}^{\exists}_{\mathsf{C}\sqcap}$ be the set $\mathsf{N}^{\exists}_{\mathsf{C}} \cup \{D \sqcap E \mid D, E \in \mathsf{N}^{\exists}_{\mathsf{C}}\}$. Assume $\Delta_{\mathcal{K}} := \{c_D \mid D \in \mathsf{N}^{\exists}_{\mathsf{C}\sqcap},\ D$ is satisfiable w.r.t. $\mathcal{K}\}$ is disjoint from N_I. Given a role S, we write $\overline{S}$ as the result of switching between a role name and its inverse: $\overline{S} = S^-$ if $S \in \mathsf{N}_\mathsf{R}$ and $\overline{S} = R$ if $S = R^-$ with $R \in \mathsf{N}_\mathsf{R}$.

Definition 1 (Canonical Model)**.** *The canonical model $\mathcal{I}_{\mathcal{K}}$ for a satisfiable DL-Lite$^{\mathcal{H}}$ KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is:*

- $\Delta^{\mathcal{I}_{\mathcal{K}}} := \mathsf{N}_\mathsf{I} \cup \Delta_{\mathcal{K}}$, $\quad a^{\mathcal{I}_{\mathcal{K}}} := a$, *for all* $a \in \mathsf{N}_\mathsf{I}$,
- $A^{\mathcal{I}_{\mathcal{K}}} := \{a \in \mathsf{N}_\mathsf{I} \mid \mathcal{K} \models A(a)\} \cup \{c_D \in \Delta_{\mathcal{K}} \mid \mathcal{K} \models D \sqsubseteq A\}$ *for all* $A \in \mathsf{N}_\mathsf{C}$,
- $R^{\mathcal{I}_{\mathcal{K}}} := \{(a, b) \in \mathsf{N}_\mathsf{I} \times \mathsf{N}_\mathsf{I} \mid \mathcal{K} \models R(a, b)\} \cup$
  $\{(a, c_{\exists S}) \in \mathsf{N}_\mathsf{I} \times \Delta_{\mathcal{K}} \mid \mathcal{K} \models \exists\overline{S}(a),\ \mathcal{K} \models \overline{S} \sqsubseteq R\} \cup$
  $\{(c_{\exists S}, a) \in \Delta_{\mathcal{K}} \times \mathsf{N}_\mathsf{I} \mid \mathcal{K} \models \exists\overline{S}(a),\ \mathcal{K} \models S \sqsubseteq R\} \cup$
  $\{(c_{\exists S}, c_{\exists\overline{S}}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} \mid \mathcal{K} \models S \sqsubseteq R\} \cup$
  $\{(c_D, c_{\exists S}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} | \mathcal{K} \models D \sqsubseteq \exists\overline{S},\ \mathcal{K} \models \overline{S} \sqsubseteq R\} \cup$
  $\{(c_{\exists S}, c_D) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} \mid \mathcal{K} \models D \sqsubseteq \exists\overline{S},\ \mathcal{K} \models S \sqsubseteq R\}$,
  for all $R \in \mathsf{N}_\mathsf{R}$.

Theorem 1. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB and let $\mathcal{I}_{\mathcal{K}}$ be the canonical model of $\mathcal{K}$. Then, for all DL-Lite$^{\mathcal{H}}$ axioms α, we have that $\mathcal{I}_{\mathcal{K}} \models \alpha$ iff $\mathcal{K} \models \alpha$.*

3 BoxLitE Semantics

Here we introduce a semantics for DL-Lite$^{\mathcal{H}}$ inspired by box-based geometric models (Abboud et al., 2020) and suitable for defining a KB embedding model that can be optimized using convex optimization. We define a geometric model that uses axis-aligned hyper-rectangles, called *boxes*.

Let ϵ and s_Ω be fixed but arbitrary [3] positive constants with $\epsilon \leq s_\Omega$. Let $\boldsymbol{\epsilon}$ and $\mathbf{s_\Omega}$ denote the d-dimensional vectors whose value in each dimension is equal to ϵ and s_Ω respectively. Henceforth, we denote elementwise comparison operators with $\leq_d$ and $\geq_d$. Using $\boldsymbol{\epsilon}$ and $\mathbf{s_\Omega}$, we define an axis-aligned hyper-rectangle, called the *universe box*:

$$\Omega = \{\mathbf{x} \in \mathbb{R}^d \mid -\mathbf{s_\Omega} \leq_d \mathbf{x} \leq_d \mathbf{s_\Omega}\}$$

in our d-dimensional Euclidean space, where $d > 0$. The universe box is useful to establish basic properties of boxes (Theorem 2) that are needed in our faithfulness proofs.

Definition 2. *A* box *is an element of the set* Box*, defined as:*

$$\mathsf{Box} := \{\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U} - \boldsymbol{\epsilon}\} \mid \mathbf{0} \leq_d (\mathbf{U} - \mathbf{L}) \leq_d 2\mathbf{s_\Omega},\ \mathbf{L}, \mathbf{U} \in \mathbb{R}^d\}.$$

For a non-empty box, the lower *and* upper *bounds are denoted* $\mathbf{L}$ *and* $\mathbf{U}$ *(resp.) of* $\mathbf{X}$*. If* $\mathbf{X}$ *is the empty box then we set* $\mathbf{L} := \mathbf{U} := \mathbf{0}$*. We denote with* $\mathbf{x}[i]$ *the* i*-th dimension of a vector* $\mathbf{x}$*. If* $\mathbf{X}$ *is a box with lower and upper bounds* $\mathbf{L}$ *and* $\mathbf{U}$*, resp., then* $\mathbf{X}[i]$ *is the pair* $(\mathbf{L}[i], \mathbf{U}[i])$.

Box Definition Intuition. The definition of Box allows boxes to have a width of $\leq 2\mathbf{s_\Omega}$ and to span outside of the universe Ω, while we require vectors associated with individual names to be within Ω (as we will see in Definition 3). This design is motivated (in the spirit of Abboud et al. (2020)) by associating any individual name with a position and a bump vector, where the embeddings of a pair of individual names (a, b) is retrieved by translating ("*bumping*") the position of a with the bump of b and vice versa. Thus, if both the position and bump vectors are within the universe box, i.e., bounded by $\mathbf{s_\Omega}$, then the embeddings of any pair of individuals is bounded by $2\mathbf{s_\Omega}$. As we associate any concept and role name with a set of boxes that shall contain the embeddings of individual tuples, it is sufficient that the widths of boxes are bounded by $2\mathbf{s_\Omega}$. The intuition for $\boldsymbol{\epsilon}$ is that most convex optimization tools do not allow strict inequalities or handle them less efficiently. However, we need to be able to define empty boxes for concepts/roles that are unsatisfiable. So we employ in Definition 2 *nonstrict inequalities* and add a small positive $\boldsymbol{\epsilon}$ to emulate strict ones.

Definition 3. *A* box interpretation η *is a function that maps:*

- *each individual name* $a \in \mathsf{N}_\mathsf{I}$ *to two vectors* $\eta(a) = (\mathsf{pos}(a), \mathsf{bump}(a))$*, namely, a position* $\mathsf{pos}(e) \in \Omega$ *and a bump* $\mathsf{bump}(e) \in \Omega$*;*
- *each concept name* $A \in \mathsf{N}_\mathsf{C}$ *to a box* $\eta(A) \in \mathsf{Box}$*;*
- *each role name* $R \in \mathsf{N}_\mathsf{R}$ *to three boxes* $\eta(R) = (\mathsf{Head}(R), \mathsf{Tail}(R), \mathsf{Bump}(R))$*, which we call* R*'s head* $\mathsf{Head}(R)$*, tail* $\mathsf{Tail}(R)$ *and bump box* $\mathsf{Bump}(R)$.

We extend the mapping function η *to arbitrary DL-Lite*$^{\mathcal{H}}$ *concept and role expressions as follows:*

$$\eta(\neg C) := \overline{\eta(C)}, \quad \eta(R^-) := (\mathsf{Tail}(R), \mathsf{Head}(R), \mathsf{Bump}(R)),$$

$$\eta(\exists R) := \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L^H_R} - \mathbf{U^B_R} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U^H_R} - \mathbf{L^B_R} - \boldsymbol{\epsilon}\}$$

where $\mathbf{L}^X_\mathbf{R}$*,* $\mathbf{U}^X_\mathbf{R}$ *with* $X \in \{\mathbf{H}, \mathbf{T}, \mathbf{B}\}$ *are the lower and upper bounds of* $\mathsf{Head}(R)$*,* $\mathsf{Tail}(R)$*, and* $\mathsf{Bump}(R)$*; and where* $\overline{\eta(C)}$ *represents* $\eta(C)$*'s complement box. Furthermore, for any* $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$*, the complement box* $\overline{\eta(C)}$ *is defined as follows:*

$$\overline{\eta(C)} := \{\mathbf{x} \in \mathbb{R}^d \mid (-\mathbf{s_\Omega} - \mathbf{L_C} + \boldsymbol{\epsilon}) \leq_d \mathbf{x} \leq_d (\mathbf{s_\Omega} - \mathbf{U_C} - \boldsymbol{\epsilon})\}$$

with $\mathbf{L_C}$ *and* $\mathbf{U_C}$ *being* $\eta(C)$*'s lower and upper bounds.*

Box Interpretation Intuition. At first sight, an intuitive definition for the complement of a box would be the set complement. However, we cannot use this notion as it leads to non-convexity. Thus, a different convex-preserving definition of box complements is required that satisfies important properties of the usual complement, as shown in Theorem 2. Next, the box interpretation of inverse roles $\eta(R^-)$ swaps the head and tail boxes of $\eta(R)$. As for the existential boxes, we define $\eta(\exists R)$ as $\mathsf{Head}(R)$ enlarged by the boundaries of

[3] In our work and, in particular, in the implementation, we take $0 < \epsilon < \epsilon_{max}$, where $\epsilon_{max} = 0.5$ and $\epsilon_{max} \leq s_\Omega/8$. Making ϵ and s_Ω learnable parameters would not enhance the representation capabilities of our model but allow for infinitely many equivalent solutions of our learned embeddings which only differ in scale.

$\mathsf{Bump}(R)$, to ensure that it contains the embeddings of all subjects of R assertions, following $\exists R$'s semantics.

Theorem 2. *For any box interpretation η:*

i) *for all $C \in \mathsf{N}_\mathsf{C}^\exists$, $\overline{\eta(C)} \in \mathsf{Box}$;*

ii) *for all $C \in \mathsf{N}_\mathsf{C}^\exists$, $\overline{\overline{\eta(C)}} = \eta(C)$;*

iii) *for all $C, D \in \mathsf{N}_\mathsf{C}^\exists$, if $\eta(C) \subseteq \eta(D)$ then $\overline{\eta(D)} \subseteq \overline{\eta(C)}$.*

Box Consistency. A box interpretation is *box consistent* if for all $C \in \mathsf{N}_\mathsf{C}^\exists$ we have that $\eta(C) \cap \overline{\eta(C)} = \emptyset$.

Definition 4. *A box interpretation η satisfies*

- *$R(a, b)$ with $R \in \mathsf{N}_\mathsf{R}$ and $a, b \in \mathsf{N}_\mathsf{I}$ iff*
$$i)\ \mathsf{pos}(a) + \mathsf{bump}(b) \in \mathsf{Head}(R)$$
$$ii)\ \mathsf{pos}(b) + \mathsf{bump}(a) \in \mathsf{Tail}(R)$$
$$iii)\ \mathsf{bump}(a), \mathsf{bump}(b) \in \mathsf{Bump}(R)$$
- *$R \sqsubseteq S$ with $R, S \in \mathsf{N}_\mathsf{R}^-$ iff* $\mathsf{Bump}(R) \subseteq \mathsf{Bump}(S)$
$$\mathsf{Head}(R) \subseteq \mathsf{Head}(S) \qquad \mathsf{Tail}(R) \subseteq \mathsf{Tail}(S)$$
- *$C \sqsubseteq D$, $C \in \mathsf{N}_\mathsf{C}^\exists$ and $D \in \mathsf{N}_{\mathsf{C}\neg}^\exists$ iff* $\eta(C) \subseteq \eta(D)$.

Regarding concept assertions with $C \in \mathsf{N}_\mathsf{C}^\exists$ and $a \in \mathsf{N}_\mathsf{I}$, η satisfies $C(a)$ if $\mathsf{pos}(a) \in \eta(C)$ *and η falsifies $C(a)$ if* $\mathsf{pos}(a) \in \overline{\eta(C)}$. *Otherwise the status of $C(a)$ is* unknown[4].

We write $\eta \models \alpha$ to indicate that η satisfies an axiom α, and write $\eta \not\models \alpha$ otherwise. Also, if $\eta \models \alpha$ for every axiom α in a KB $\mathcal{K}$ then we say that η satisfies $\mathcal{K}$ or, equivalently, that η is a model of $\mathcal{K}$, in symbols, $\eta \models \mathcal{K}$.

The most intricate part of Definition 4 corresponds to role assertions, so we explain this in more details. Items $i)$ and $ii)$ are as in Abboud et al. (2020). Item $iii)$ is helpful to embed existential concepts. Recall that $\eta(\exists R)$ is $\mathsf{Head}(R)$ enlarged by the boundaries of $\mathsf{Bump}(R)$. Thus, if $R(a, b)$ is satisfied in a box interpretation, $\mathsf{pos}(a)$ may lie at most the size of $\mathsf{Bump}(R)$ away from $\mathsf{Head}(R)$.

With this, we have defined BoxLitE's semantics in terms of box interpretations. In contrast to (Abboud et al., 2020), we (i) define *complement boxes*, which allow the formulation of convex constraints that guarantee the satisfaction of concept disjointness axioms in our embedding solutions and (ii) associate roles with additional *bump* boxes that constrain the maximal bump of an individual that is accepted by a role embedding. The latter is important to distinguish between axioms of the form $R \sqsubseteq S$ and $\{\exists R \sqsubseteq \exists S, \exists R^- \sqsubseteq \exists S^-\}$. In the next section, we will define faithfulness properties of geometric models. Furthermore, we will analyse which faithfulness properties for DL-Lite$^\mathcal{H}$ can be satisfied by BoxLitE's box interpretations.

[4]Geometric models for languages that have negation normally need to allow for the 'unknown' truth status. This happens in the Cone Semantics (Lütfü Özçep, Leemhuis, and Wolter, 2020), tailored for $\mathcal{ALC}$, which has negation. Since DL-Lite$^\mathcal{H}$ allows for $B \sqsubseteq \neg C$ we use complement boxes and the 'unknown' truth state.

4 Model Faithfulness

In this section, we study model faithfulness (Lütfü Özçep, Leemhuis, and Wolter, 2020), which is a property that is useful to show that geometric models correctly represent the knowledge present in KBs. For KB completion, one is particularly interested in preserving the conceptual knowledge in TBoxes while allowing new assertions to hold.

Definition 5. *(Adapted (Lütfü Özçep, Leemhuis, and Wolter, 2020; Bourgaux et al., 2024)) Let $\mathcal{L}$ be an ontology language and $\mathcal{K}$ a satisfiable KB in $\mathcal{L}$. A box interpretation η is*

- weakly KB faithful *for $\mathcal{L}$ and $\mathcal{K}$ if for every KB axiom α in $\mathcal{L}$, $\eta \models \alpha$ implies that α is consistent with $\mathcal{K}$;*
- KB entailed *for $\mathcal{L}$ and $\mathcal{K}$ if for every KB axiom α in $\mathcal{L}$ that is entailed by $\mathcal{K}$, $\eta \models \alpha$.*

*We may omit "weakly" and just say "faithful". We may also omit $\mathcal{L}$ and/or $\mathcal{K}$ if they are clear from the context. If a box η interpretation satisfies both of the properties we say that η is a KB **faithful model**. These notions can be adapted for the case in which $\mathcal{L}$ is a description logic with TBox and ABox axioms and for the case where we consider only TBox axioms (or only ABox axioms) for TBox (or ABox) faithful models.*

Proposition 1 and Proposition 2 give basic conditions for faithfulness in DL-Lite$^\mathcal{H}$.

Proposition 1. *Let $\mathcal{T}$ be a DL-Lite$^\mathcal{H}$ TBox. If $\eta \models \mathcal{T}$ then η is (weakly) TBox faithful.*

Proposition 1 follows from the fact that, in DL-Lite$^\mathcal{H}$, TBoxes are always consistent, that is, inconsistencies only happen when considering a TBox and an ABox.

Proposition 2. *Let $\mathcal{K}$ be a DL-Lite$^\mathcal{H}$ KB and η a box interpretation that is box consistent. If $\eta \models \mathcal{K}$ then η is (weakly) KB faithful.*

We are now ready to state Theorem 3, which is the main result of this section.

Theorem 3. *There exists a suitable s_Ω such that every satisfiable DL-Lite$^\mathcal{H}$ KB $\mathcal{K}$ has a box interpretation $\eta_{\mathcal{I}_\mathcal{K}}$ that is a KB faithful model and box consistent.*

Sketch. The proof strategy of Theorem 3 consists of first defining a mapping that translates classical finite interpretations into box interpretations. Then, given a DL-Lite$^\mathcal{H}$ KB $\mathcal{K}$, we construct the canonical model $\mathcal{I}_\mathcal{K}$ in Definition 1 and use this mapping to create a box interpretation $\eta_{\mathcal{I}_\mathcal{K}}$ (using $s_\Omega = 4$) that is a KB faithful model of $\mathcal{K}$. In fact, our proof provides a stronger guarantee: the embedding is KB entailed and strongly KB faithful (Lütfü Özçep, Leemhuis, and Wolter, 2020; Bourgaux et al., 2024). ☐

Informally, Theorem 3's proof proposes for every satisfiable DL-Lite$^\mathcal{H}$ KB $\mathcal{K}$ an algorithm for constructing a box interpretation (i.e., a BoxLitE embedding) that is KB faithful and box consistent. Using this algorithm, we can compute an upper bound for the minimal number of dimensions that a box interpretation requires to satisfy certain faithfulness properties, leading to Corollaries 1 and 2.

Corollary 1. *Let $\mathcal{K}$ be a DL-Lite$^{\mathcal{H}}$ KB with an empty ABox. Then for any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that $d_{\min} = |\mathsf{N_C}| + 3|\mathsf{N_R}|$ and η is a TBox faithful model of $\mathcal{K}$.*

Corollary 2. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB. Then for any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that: $d_{\min} = |\mathsf{N_C}| + |\mathsf{N_R}|(2 + |\mathsf{N_I}| + 2|\mathsf{N_R}|)$ and η is a KB faithful model of $\mathcal{K}$.*

With this we have finished the theoretical analysis of BoxLitE's faithfulness properties. What now remains to show is how our BoxLitE approach can be translated into a convex optimization problem that *ensures* certain faithfulness properties. We will investigate this in the next section.

5 BoxLitE's Convex Optimization Problem

Here, we formulate the representation of KBs with BoxLitE as a constrained convex optimization problem (Boyd and Vandenberghe, 2004; Rockafellar, 1997) and investigate how the formulation relates to faithfulness properties. In Section 5.1, we describe how TBox axioms are translated to constraints. In Section 5.2, we discuss the objective and scoring functions. In Section 5.3, we establish that our problem formulation ensures the KB faithful and TBox entailed properties for DL-Lite$^{\mathcal{H}}$. In summary, the TBox corresponds to a part of the constraints of the convex optimization problem and the ABox corresponds to a part of the objective function of the problem. In this way, a solver can handle both the ABox and TBox simultaneously by minimizing the objective function subject to the constraints.

5.1 From TBox Axioms to Constraints

Given a box interpretation η, we concatenate all of η's parameters into a single vector z of $n := (2|\mathsf{N_I}| + 2|\mathsf{N_C}| + 6|\mathsf{N_R}|)d$ dimensions[5]. We call this vector $z \in \mathbb{R}^n$ an *embedding solution*. Conversely, given such a $z \in \mathbb{R}^n$, we can reconstruct η. However, not all z will lead to an η that satisfies Definition 3. Therefore, we devise constraints that z must satisfy, such that the corresponding η has desirable properties.

In particular, given an arbitrary DL-Lite$^{\mathcal{H}}$ TBox $\mathcal{T}$, we want to ensure that any feasible solution for a box interpretation (i) is box consistent, i.e., for all $C \in \mathsf{N}_\mathsf{C}^{\exists}$ it holds that $\eta(C) \cap \eta(\neg C) \neq \emptyset$, (ii) ensures that any box width is positive and bounded by $2\mathbf{s_\Omega}$, (iii) ensures that any individual embedding is within the universe box, and (iv) guarantees that any TBox axiom is satisfied in the embedding solution. In the following, we define convex constraints for Points (i)–(iv).

Box Consistency Constraints. Intuitively, Point (i) ensures that the set of feasible solutions only contains those solutions that do not predict any contradictions. However, enforcing box consistency directly is non-convex, due to the need of computing intersection boxes (see Section 7). Thus, we need to define a convex alternative that guarantees non-overlap of $\eta(C)$ and $\eta(\neg C)$. We ensure Point (i) by reserving one dimension i_C per $C \in \mathsf{N}_\mathsf{C}^{\exists}$ and defining one constraint per dimension i_C, formally described in (1). By reserving one dimension i_C per $C \in \mathsf{N}_\mathsf{C}^{\exists}$, the minimal number of dimensions of our method depends on the TBox. In particular, our method requires at least $|\mathsf{N_C}| + 2|\mathsf{N_R}|$ dimensions [6] for a TBox with $|\mathsf{N_C}|$ concept names and $|\mathsf{N_R}|$ role names.

$$\frac{\mathbf{L_C}[i_C] + \mathbf{U_C}[i_C]}{2} \leq -\frac{s_\Omega}{2} \tag{1}$$

The inequality in (1) ensures box consistency (see Theorem 6 in the appendix).

Box Width Constraints. As mentioned in Point (ii), our method assumes that the width of any box with lower and upper bounds $\mathbf{L}$ and $\mathbf{U}$ is positive and bounded by $2\mathbf{s_\Omega}$, i.e., $\mathbf{0} \leq_d \mathbf{U} - \mathbf{L} \leq_d 2\mathbf{s_\Omega}$. The following inequalities ensure this requirement:

$$\begin{aligned} \mathbf{0} &\leq_d \mathbf{U_C} - \mathbf{L_C} \leq_d 2\mathbf{s_\Omega} \\ \mathbf{0} &\leq_d \mathbf{U}_\mathbf{R}^X - \mathbf{L}_\mathbf{R}^X \leq_d 2\mathbf{s_\Omega}, \end{aligned} \tag{2}$$

where $C \in \mathsf{N_C}$, $R \in \mathsf{N_R}$, and $X \in \{\mathbf{H}, \mathbf{T}, \mathbf{B}\}$.

Universe Constraints. As mentioned in Point (iii), our method assumes that positions and bumps of individual embeddings are within the universe box Ω (see Definition 3). The following inequalities ensure this property:

$$\begin{aligned} -\mathbf{s_\Omega} + \boldsymbol{\epsilon} &\leq_d \mathsf{pos}(a) \leq_d \mathbf{s_\Omega} - \boldsymbol{\epsilon} \\ -\mathbf{s_\Omega} + \boldsymbol{\epsilon} &\leq_d \mathsf{bump}(a) \leq_d \mathbf{s_\Omega} - \boldsymbol{\epsilon}, \end{aligned} \tag{3}$$

where $a \in \mathsf{N_I}$.

TBox Axiom Constraints. Finally, to guarantee the satisfaction of TBox axioms in feasible solutions (Point (iv)), first we need to define the boundaries of concept and role embeddings. Based on the definitions in Section 3, the boundaries of $\eta(\neg C)$, $\eta(\exists R)$, and $\eta(R^-)$ are defined as follows:

$$\begin{aligned} \mathbf{L_{\neg C}} &:= -\mathbf{s_\Omega} - \mathbf{L_C},\ \mathbf{U_{\neg C}} := \mathbf{s_\Omega} - \mathbf{U_C}, \\ \mathbf{L_{\exists R}} &:= \mathbf{L_R^H} - \mathbf{U_R^B},\ \mathbf{U_{\exists R}} := \mathbf{U_R^H} - \mathbf{L_R^B}, \\ \mathsf{Head}(R^-) &:= \mathsf{Tail}(R),\ \mathsf{Tail}(R^-) := \mathsf{Head}(R), \\ \mathsf{Bump}(R^-) &:= \mathsf{Bump}(R). \end{aligned}$$

Let $B \in \mathsf{N}_\mathsf{C}^{\exists}$, $C \in \mathsf{N}_{\mathsf{C}_\neg}^{\exists}$, and $R, S \in \mathsf{N}_\mathsf{R}^-$. To satisfy concept and role inclusions within any embedding solution z (Definition 4), we employ the following linear inequalities:

- $B \sqsubseteq C$ corresponds to $\mathbf{L_C} \leq_d \mathbf{L_B}$ and $\mathbf{U_B} \leq_d \mathbf{U_C}$,
- $R \sqsubseteq S$ corresponds to $\mathbf{L}_\mathbf{S}^X \leq \mathbf{L}_\mathbf{R}^X, \mathbf{U}_\mathbf{R}^X \leq \mathbf{U}_\mathbf{S}^X$ with $X \in \{\mathbf{H}, \mathbf{T}, \mathbf{B}\}$.

Problem Size of the Translation. We point out that the size of BoxLitE's problem formulation increases linearly w.r.t. the size of the KB $\mathcal{K}$ and $|\mathsf{N_C} \cup \mathsf{N_R} \cup \mathsf{N_I}|$. Let $\mathcal{T}$ be $\mathcal{K}$'s TBox, the number of inequalities, as described in this section, gives us $O((|\mathsf{N_I}| + |\mathsf{N_C}| + |\mathsf{N_R}| + |\mathcal{T}|)d)$ constraints.

[5]For each individual in $\mathsf{N_I}$ we have two vectors, one for the position and one for the bump. For each concept in $\mathsf{N_C}$ we also have two vectors, one for the upper corner and one for the lower corner of the box. For each role in $\mathsf{N_R}$ we have 3 boxes: head, tail, and bump; each box needs the upper and lower corner vectors.

[6]We have $2|\mathsf{N_R}|$ because of $\exists R$ and $\exists R^-$ for each $R \in \mathsf{N_R}$.

5.2 Optimization: Ranking Assertions

When ranking assertions, we would like the distance between the position of an individual w.r.t to a box associated with a concept name to affect the score of the corresponding assertion: the smaller the distance to the elements of the box, the higher the score. The same intuition also applies for roles, but needs to take into account the way role assertions are satisfied in the embedding model (Definition 4).

To express this intuition, we employ a *signed distance* function. The signed distance assigns nonpositive values to the assertions satisfied in the embedding, and positive values to those that are violated. Moreover, it reflects *how* an assertion is satisfied (or violated): for concept assertions, individuals with embeddings deeper inside the box have a smaller negative value than those closer to the border, while individuals with embeddings outside the box have positive values, and those values are higher the farther away the embeddings of the individuals are from the box.

Signed Distance. We move on to the formal definition of signed distance to a set. First, the *Euclidean distance to a set* $S \subseteq \mathbb{R}^n$ is defined as the function $\mathsf{dist}_e : \mathbb{R}^n \to \mathbb{R}$ such that $\mathsf{dist}_e(y, S) := \inf\{\|x - y\|_2 \mid x \in S\}$, $\forall y \in \mathbb{R}^n$. Let S^c be $\mathbb{R}^n \setminus S$. Then, the *signed distance* (also called *oriented distance*) to S is

$$\mathsf{sdist}(y, S) := \begin{cases} \mathsf{dist}_e(y, S) & \text{if } y \notin S \\ -\mathsf{dist}_e(y, S^c) & \text{if } y \in S, \end{cases} \tag{4}$$

e.g., see Chapter 7 of Delfour and Zolésio (2011). Given $y \in \mathbb{R}^n$, we denote by $y^+ \in \mathbb{R}^n$ the vector that corresponds to replacing the negative components of y by 0. E.g., $(1, -2, 3, -4)^+ = (1, 0, 3, 0)$.

Proposition 3. *The signed distance to* $\mathbb{R}^n_-$ *satisfies*

$$\mathsf{sdist}(y, \mathbb{R}^n_-) = \begin{cases} \|y^+\|_2 & \text{if } y \notin \mathbb{R}^n_- \\ \max_{i \in \{1,\dots,n\}} y_i & \text{if } y \in \mathbb{R}^n_-, \end{cases} \tag{5}$$

where $\mathbb{R}^n_- = \{y \in \mathbb{R}^n \mid y_i \leq 0, 1 \leq i \leq n\}$.

Let $\mathbf{B}$ be a box with bounds $\mathbf{L}$ and $\mathbf{U}$. We define $\mathsf{dist}(\mathbf{B}, \mathbf{x})$ as the signed distance of the concatenated vector $(\mathbf{L} + \boldsymbol{\epsilon} - \mathbf{x}) \oplus (\mathbf{x} - \mathbf{U} + \boldsymbol{\epsilon})$ to $\mathbb{R}^{2d}_-$. That is,

$$\mathsf{dist}(\mathbf{B}, \mathbf{x}) := \mathsf{sdist}((\mathbf{L} + \boldsymbol{\epsilon} - \mathbf{x}) \oplus (\mathbf{x} - \mathbf{U} + \boldsymbol{\epsilon}), \mathbb{R}^{2d}_-).$$

We now define the loss terms using dist.

Concept Assertion Loss. We define the loss for concept assertions $D(a)$ as follows:

$$\mathcal{L}_{\mathsf{concept}}(D, a) := \mathsf{dist}(\eta(D), \mathsf{pos}(a)).$$

Intuitively, minimizing the loss $\mathcal{L}_{\mathsf{concept}}(D, a)$ of a concept assertion $D(a)$ pushes an individual's position embedding $\mathsf{pos}(a)$ into D's concept embedding $\eta(D)$.

Role Assertion Loss. Similarly, the loss $\mathcal{L}_{\mathsf{role}}(S, a, b)$ of a role assertion $S(a, b)$ pushes the translated individual embedding $\mathsf{pos}(a) + \mathsf{bump}(b)$ and, respectively, $\mathsf{pos}(b) + \mathsf{bump}(a)$ into S's head box $\mathsf{Head}(S)$ and S's tail box $\mathsf{Tail}(S)$. Furthermore, it pushes the bumps of a and b into S's bump box $\mathsf{Bump}(S)$. Based on the distance function dist, we define the loss for role assertions $S(a, b)$ as follows:

$$\begin{aligned} \mathcal{L}_{\mathsf{role}}(S, a, b) := \max\big(&\mathsf{dist}(\mathsf{Head}(S), \mathsf{pos}(a) + \mathsf{bump}(b)), \\ &\mathsf{dist}(\mathsf{Tail}(S), \mathsf{pos}(b) + \mathsf{bump}(a)), \\ &\mathsf{dist}(\mathsf{Bump}(S), \mathsf{bump}(a)), \\ &\mathsf{dist}(\mathsf{Bump}(S), \mathsf{bump}(b))\big). \end{aligned}$$

In practice, KBs typically contain only positive assertions and no explicit negative ones. Therefore, if we were to minimise only the loss terms of the assertions, BoxLitE would tend to assign high scores to all potential assertions, not distinguishing between true and false assertions. A common way to address this issue in KB embedding methods is *negative sampling* (Bordes et al., 2013; Sun et al., 2019; Abboud et al., 2020). In that approach, given a role assertion $R(h, t)$ contained in the KB, one constructs a *corrupted* assertion by exchanging either the head h or tail t of the role assertion by a randomly chosen individual from $\mathsf{N_I}$. Since the number of assertions that hold is usually much smaller than the number of all possible assertions that can be made, most corrupted assertions tend to be false. KB embedding models are then trained to assign high scores to ABox assertions in the training data and low scores to corrupted ones. However, negative sampling leads to nonconvex loss terms, which are often incompatible with convex optimization (see Section 7); and it is computationally expensive, as competitive performance often requires sampling hundreds or even thousands of negative assertions per ABox assertion (Abboud et al., 2020; Lu and Hu, 2020; Pavlović and Sallinger, 2023b). Next, we adapt negative sampling to concepts in $\mathsf{N}^{\exists}_{\mathsf{C}}$, leading to convex regularization terms that are fast to compute and push individual embeddings into complement boxes as required.

Negative Concept Regularization. Our approach builds on three observations: (1) directly pushing individual embeddings outside a box (i.e., into its geometric complement) leads to nonconvex optimization terms; (2) BoxLitE introduces convex boxes representing the complement of concept embeddings, which we can use to regularize the scores; and (3) BoxLitE does not offer convex representations for the complement of role embeddings, yet a role's domain and range can be expressed as existential concept embeddings. Based on these observations, we introduce a negative concept regularization term that for any potential concept assertion $D(a) \notin \mathcal{A}$ with $D \in \mathsf{N}^{\exists}_{\mathsf{C}}$ and $a \in \mathsf{N_I}$ pushes $\mathsf{pos}(a)$ into $\overline{\eta(D)}$, reducing the score of $D(a)$. This regularization term keeps plausibility scores for arbitrary assertions low, while the assertion loss terms selectively increase the scores of ABox assertions explicitly contained in the KB:

$$\mathcal{L}_{\mathsf{negative}}(D, a) := \mathcal{L}_{\mathsf{concept}}(\neg D, a). \tag{6}$$

Box Width Regularization. A second strategy to keep scores for arbitrary assertions within a reasonable range is to regularize the box size of concept and role embeddings. Specifically, for any concept name $D \in \mathsf{N_C}$ we regularize the size of its concept box $\eta(D)$; and for any role name $S \in \mathsf{N_R}$, we regularize the size of its head $\mathsf{Head}(S)$, tail $\mathsf{Tail}(S)$, and bump box $\mathsf{Bump}(S)$. Formally, we define the box regularization term $\mathcal{R}_{\mathsf{width}}(\mathbf{B})$ of a box $\mathbf{B}$ with bounds $\mathbf{U}$ and $\mathbf{L}$, as:

$$\mathcal{R}_{\mathsf{width}}(\mathbf{B}) := \|\mathbf{U} - \mathbf{L}\|_2. \quad (7)$$

Objective Function. We now assemble these loss and regularization terms into our objective function:

$$\max\left(\max_{D(a)\in\mathcal{A}} \mathcal{L}_{\mathsf{concept}}(D,a), \max_{S(a,b)\in\mathcal{A}} \mathcal{L}_{\mathsf{role}}(S,a,b)\right) + \lambda_1 \max_{D(a)\in N_C^{\exists}\setminus\mathcal{A}} \mathcal{L}_{\mathsf{negative}}(D,a) + \lambda_2\Big(\sum_{D\in\mathsf{N_C}} \mathcal{R}_{\mathsf{width}}(\eta(D)) + \sum_{S\in\mathsf{N_R}} \mathcal{R}_{\mathsf{width}}(\mathsf{Head}(S)) + \mathcal{R}_{\mathsf{width}}(\mathsf{Tail}(S))\Big) + \lambda_3 \sum_{S\in\mathsf{N_R}} \mathcal{R}_{\mathsf{width}}(\mathsf{Bump}(S)).$$

Scores. To rank the assertions, we define two scoring functions, one for concept and one for role assertions.

The score $s(D,a)$ of concept assertions $D(a)$ is:

$$s(D,a) := -\,\mathcal{L}_{\mathsf{concept}}(D,a)$$

and the score $s(S,a,b)$ of role assertions $S(a,b)$ is:

$$s(S,a,b) := -\,\mathcal{L}_{\mathsf{role}}(S,a,b).$$

The scoring functions of the concept and role assertions are the negative of the corresponding loss functions. The intuition for this is that solving the optimization problem, i.e., minimizing the concept and role assertion loss for ABox assertions, maximizes their score.

5.3 DL-Lite$^{\mathcal{H}}$ KB Faithfulness

Translating TBox axioms to linear inequalities that are used as convex constraints in the optimization problem (see Section 5.1), guarantees any BoxLitE embedding solution z to satisfy the concept inclusions in the TBox and to be KB faithful for DL-Lite$^{\mathcal{H}}$.

Let $\mathcal{C}_{\mathcal{K}} \subseteq \mathbb{R}^n$ be the set of z's such that the constraints of Section 5.1 are satisfied. That is, $z \in \mathcal{C}_{\mathcal{K}}$ if and only if: (a) for each TBox axiom the corresponding inequalities are satisfied; and (b) the box consistency and universe constraints are satisfied, see the appendix for details. Also, let $f_{\boldsymbol{\lambda}} : \mathbb{R}^n \to \mathbb{R}$ be the function that maps z to the objective value for a given choice of nonnegative hyperparameters $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3_+$.

Theorem 4. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB. For nonnegative $\boldsymbol{\lambda}$ the following optimization problem is convex.*

$$\min_{z\in\mathbb{R}^n} f_{\boldsymbol{\lambda}}(z), \quad \text{subject to } z \in \mathcal{C}_{\mathcal{K}}. \quad (8)$$

In particular, $f_{\boldsymbol{\lambda}}$ is a convex function, $\mathcal{C}_{\mathcal{K}}$ is a polyhedral set and the following items hold.

i) For d as in Corollary 1, and s_Ω as in Theorem 3, $\mathcal{C}_{\mathcal{K}}$ is nonempty.

ii) Any embedding solution $z \in \mathcal{C}_{\mathcal{K}}$ corresponds to a box consistent interpretation that is TBox faithful.

iii) If $\lambda_1 = 0$ and there is an optimal solution z^* such that $f_{\boldsymbol{\lambda}}(z^*) \le 0$ then the box interpretation corresponding to z^* is KB faithful.

iv) Suppose that $\lambda_1 = \lambda_2 = \lambda_3 = 0$. For d as in Corollary 2, and s_Ω as in Theorem 3, there is an optimal solution z^* s.t. $f_{\boldsymbol{\lambda}}(z^*) \le 0$ holds.

Informally, Theorem 4 states that we can find a box interpretation for a satisfiable $\mathcal{K}$ via convex optimization over a polyhedral set. Any $z \in \mathcal{C}_{\mathcal{K}}$ (whether optimal or not) corresponds to a box interpretation η that is TBox faithful and box consistent. Item $i)$ gives a bound on the minimum required d to ensure that $\mathcal{C}_{\mathcal{K}}$ is nonempty, but this estimate seems to be conservative. Items $iii)$ and $iv)$ imply that if we wish to find a solution that is KB faithful, this can be done by setting the hyperparameters associated to regularization terms to 0, setting d to be sufficiently large and solving (8).

Finally, it turns out that (8) can be formulated as a *second-order cone program* (SOCP) (Lobo et al., 1998), (Ben-Tal and Nemirovski, 2001, Lecture 3).

Theorem 5. *For nonnegative $\boldsymbol{\lambda}$, the problem in* (8) *can be reformulated as an equivalent SOCP.*

Theorem 5 is important because it shows that (8) can be solved efficiently via high-quality open-source solvers such as SeDuMi (Sturm, 1999) and SDPT3 (Tütüncü, Toh, and Todd, 2003) or commercial solvers such as Gurobi (Gurobi Optimization, LLC, 2023) and MOSEK (ApS, 2026). The conversion of a problem as in (8) to a SOCP that can be handled by the aforementioned solvers, although tedious, can be automated by modelling tools for convex optimization such as CVXPY (Agrawal et al., 2018).

6 Proof of Concept

Here, we present empirical evidence for the theoretical foundations established so far. We evaluate BoxLitE's performance on subsets of the Family dataset (Imenes, Guimarães, and Ozaki, 2023), providing first results for its scalability and reasoning capabilities. In our experiments, we only consider KBs with satisfiable concepts.

Reproducibility. We implemented BoxLitE's optimization problem in Python 3.12 using CVXPY (Diamond and Boyd, 2016; Agrawal et al., 2018) for modeling and MOSEK (ApS, 2026) for solving it. In our evaluation we include experiments with other KBE approaches using classical stochastic gradient descent (SGD). We use PyKEEN 1.11.1 (Ali et al., 2021) in these experiments. We ran each of our experiments on an Apple Mac Mini Desktop Computer with M4 Chip with 10 Core CPU and 10 Core GPU: 16GB (Shared Memory). More details can be found in the appendix and the code is available at https://github.com/AleksVap/BoxLitE. We focus on the following questions.

(Q1) What is the reasoning performance and what is the effect of the regularization terms on the results?

(Q2) How does increasing the dataset size affect the prediction performance, compilation, and solving time?

(Q3) How well does our method compare with classical embedding methods based on stochastic gradient descent?

Experimental Setup. To answer each of these questions, we have created a set of datasets (F_v1-4) of varying sizes from the family dataset (Imenes, Guimarães, and Ozaki, 2023). We derived these datasets by sampling k assertions of the family dataset's ABox with forest fire sampling (Leskovec, Kleinberg, and Faloutsos, 2005), a popular sampling technique for large graphs. Furthermore, since the family dataset solely provides role assertions in its ABox, we selected all concept inclusions in the family dataset that only include roles and extended the TBox by the disjointness axiom $\exists \mathsf{hasFather}^- \sqsubseteq \neg\exists \mathsf{hasMother}^-$. We list the TBox of the created datasets in Figure 1.

Evaluation Setup. To evaluate BoxLitE's performance, we created a set of inferred role assertions by (i) adding any role assertion that logically follows from each dataset and (ii) removing any assertion that occurs in the ABox. We randomly split this set into a validation set (20%), used for model selection, and a test set (80%), used for evaluating the performance of the selected model. We use the standard evaluation setting for KB completion [7] (Abboud et al., 2020; Pavlović and Sallinger, 2024a; Xiong et al., 2022).

Dataset Properties. Table 1 lists the number of assertions and individuals of the train, validation, and test sets of F_v1-4. We sampled each of these datasets individually. Thus, although datasets F_v1-4 gradually increase in size, they are different from each other.

$\mathsf{relative}^- \sqsubseteq \mathsf{relative}$	$\mathsf{hasSibling} \sqsubseteq \mathsf{relative}$
$\mathsf{hasChild} \sqsubseteq \mathsf{relative}$	$\mathsf{hasParent} \sqsubseteq \mathsf{relative}$
$\mathsf{hasFather} \sqsubseteq \mathsf{hasParent}$	$\mathsf{hasMother} \sqsubseteq \mathsf{hasParent}$
$\mathsf{spouse}^- \sqsubseteq \mathsf{spouse}$	$\mathsf{hasSibling}^- \sqsubseteq \mathsf{hasSibling}$
$\mathsf{spouse} \sqsubseteq \mathsf{relative}$	$\exists \mathsf{hasFather}^- \sqsubseteq \neg\exists \mathsf{hasMother}^-$

Figure 1: TBox of datasets F_v1-4.

(Q1) Performance. In Table 2 we present the link prediction results on the test set for BoxLitE with the three regularization terms that appear in the objetive function (Section 5.2). To study the effect of each of these terms, we also performed experiments in which we remove them. We denote by BoxLitEi the version of BoxLitE without the regularization term multiplied by λ_i. The removal of each of the regularization terms decreases the overall performance of the model, with the removal of the term associated with λ_2 being the one that most negatively impacts the results.

[7]The evaluation of a KB embedding model typically needs a set of true and corrupted role assertions. True role assertions $R(a, b)$ of the KB are corrupted by replacing a or b by any $c \in \mathsf{N_I}$ such that the corrupted assertion is not within the KB. The performance of KB embedding models is typically measured using the filtered versions (Bordes et al., 2013) of the *mean reciprocal rank* (MRR) and H@k, the proportion of true assertions within the predicted assertions whose rank is at maximum k.

Name	#Train	#Val	#Test	#Individuals
F_v1	300	110	440	155
F_v2	500	209	837	233
F_v3	1001	432	1731	368
F_v4	3014	1226	4908	895

Table 1: Dataset properties: Number of train, validation, and testing assertions, and individuals.

(Q2) Scalability. The prediction performance on the test set, reported in Table 2, reduces slowly with increasing dataset sizes. Regarding the time required for each instance, we recall that a problem modelled through CVXPY is first compiled and then sent to a solver of the user's choice, which in our case is MOSEK. Given a specific choice of hyperparameters $\lambda_1, \lambda_2, \lambda_3$, Table 3 displays the compilation and solving time required for obtaining an optimal solution to the problem in Theorem 4, for each of the datasets F_v1-4. In our implementation, we tested 354 hyperparameter configurations for each dataset. While changing the hyperparameters requires solving the optimization problem again, it does not require a recompilation. Overall, compilation does not take more than a couple of seconds and all solution times were less than 30 seconds, which is quite reasonable considering that the final SOCP corresponding to F_v4 has around 3.5 million variables and 2.3 million constraints, which has size comparable to some of the instances that appear in Mittelmann's benchmark of large SOCPs (Mittelmann, 2026). Even more, the compilation time in Table 3 grows linearly with the number of training axioms, while the solving time increases only sublinearly.

(Q3) Comparison. Comparing BoxLitE with other KBEs in a direct way is tricky since KBEs in the literature consider other languages and are mostly optimized using SGD. BoxLitE is the first KBE for DL-Lite$^{\mathcal{H}}$ ontologies and the first that solves link prediction via convex optimization. In the discussion, we include an argument for why it is challenging to design convex optimization approaches for languages with conjunctions, which appear in other papers. To illustrate how our approach roughly compares with classical embedding methods such as BoxE, RotatE, and ComplEx, based on SGD, we run those methods on the ABox part of F_v1-4. We see that the results of BoxLitE are better than RotatE and ComplEx, but still behind BoxE. One exception is F_v2 where our method performs better. None of the SGD methods had any rule injections to reflect the DL-Lite$^{\mathcal{H}}$ TBox axioms used in BoxLitE. This is a disadvantage for the SGD methods. On the other hand, BoxLitE has negative sampling applied only to existential assertions, which is a disadvantage for our case. A possible main reason for the performance gap is BoxLitE's hyperparameter optimization (HPO). It currently relies on grid search, which ensures full control over the explored parameter space, vital for ablations.

Dataset	Model	MRR	H@1	H@3	H@10
F_v1	BoxLitE1	.632	.435	.800	.962
	BoxLitE2	.249	.134	.284	.475
	BoxLitE3	.698	.524	.842	.984
	BoxLitE	.720	.545	.870	.979
	BoxE	.826	.719	.918	.986
	RotatE	.474	.295	.584	.805
	ComplEx	.322	.206	.359	.535
F_v2	BoxLitE1	.428	.268	.519	.726
	BoxLitE2	.144	.072	.142	.273
	BoxLitE3	.541	.342	.668	.897
	BoxLitE	.549	.352	.685	.905
	BoxE	.432	.337	.461	.604
	RotatE	.209	.147	.218	.314
	ComplEx	.341	.265	.357	.488
F_v3	BoxLitE1	.364	.239	.396	.626
	BoxLitE2	.116	.055	.108	.224
	BoxLitE3	.414	.280	.472	.681
	BoxLitE	.433	.288	.509	.708
	BoxE	.894	.827	.946	.998
	RotatE	.177	.104	.199	.298
	ComplEx	.184	.119	.205	.284
F_v4	BoxLitE1	.339	.204	.415	.592
	BoxLitE2	.059	.023	.54	.110
	BoxLitE3	.409	.272	.483	.638
	BoxLitE	.444	.249	.571	.763
	BoxE	.626	.444	.772	.929
	RotatE	.245	.129	.300	.450
	ComplEx	.134	.087	.156	.203

Table 2: Test Results on F_v1-4. Average of 3 runs for SGD methods. Standard deviation nearly 0 in all cases.

Dataset	Compilation Time	Solving Time
F_v1	0.34	11.91
F_v2	0.54	13.14
F_v3	1.05	16.78
F_v4	3.12	26.24

Table 3: Time in seconds split by dataset for BoxLitE.

Yet, grid search does not adaptively guide the hyperparameter search. By contrast, the SGD approaches, implemented in PyKEEN (Ali et al., 2021), make use of more efficient, model-based optimization techniques. Incorporating such adaptive HPO methods in future work could lead to more effective exploration of BoxLitE's hyperparameter landscape and performance gains.

7 Discussion on Optimization in KBEs

In this section, we discuss some aspects related to differentiability and primary causes for nonconvexity in previous KB embedding works. We also recall how we address each challenge in our approach.

Nondifferentiability. In our approach nondifferentiability is not an issue because in view of Theorem 5 the underlying optimization problem can be cast as a second-order cone program. Informally, the nondifferentiable part of the problem gets embedded into the conic constraints and our solver of choice (MOSEK) (ApS, 2026) can handle this kind of problem without theoretical issues.

Negative sampling. Negative sampling as described, say, in (Sun et al., 2019, Section 3.3) includes the minimization of terms of the form

$$-\log(\sigma(\gamma - d(x))) - \sum_{i=1}^{n} w_i(\log(\sigma(d(x_i) - \gamma))),$$

where σ is a sigmoid function (e.g., $1/(1 + e^{-x})$), w_i are nonnegative weights, γ is a margin parameter, the x_i's are "negative samples" and d is a distance-like function which may include, for example, a p-norm term. Generally speaking, a function of the form $-\log(\sigma(d(z) - \gamma)$ is neither convex (nor concave) nor differentiable everywhere as a function of z. This can be seen by considering the 1-dimensional case, where d is the absolute value function and z is scalar, so that we obtain the function $-\ln(\sigma(|z| - \gamma)) = \ln(1 + e^{\gamma-|z|})$, which, albeit continuous, is nonconvex and nondifferentiable at $z = 0$. In contrast, BoxLitE adopts negative sampling in a convex way (specifically by introducing the negative concept regularization terms of Section 5.2) that push individuals into complement boxes.

Nonconvex loss terms. The loss terms proposed in previous works are often constructed through operations that do not preserve convexity in general. We briefly take a look at this issue in two works that are more closely related to our approach. To be fair, none of the works described below contain claims that their optimization problems are convex. For BoxE, it is not clear whether the distance function considered in (Abboud et al., 2020, Section 4) is convex as a *function of the parameters that need to be optimized during learning*, since it includes division and multiplication by a width term that depends on the size of the boxes. Division and multiplication, generally speaking, are not operations that preserve convexity. In BoxEL, the authors consider loss terms that are quotients of volumes of boxes (or approximations thereof), e.g., see (Xiong et al., 2022, Section 4.4). Again, division does not preserve convexity in general. In contrast, all of BoxLitE's loss terms together with the distance function in our approach are convex.

The logic fragment. A final source of nonconvexity for certain approaches seems to be the choice of the description logic fragment itself. Approaches based on minimizing loss terms together with logical languages that include, say, conjunction on the left-hand side typically lead to nonconvexity. Suppose that the conjunction of concepts C, D is interpreted as the set intersection of the corresponding boxes $\eta(C), \eta(D)$, as observed in the works for the $\mathcal{EL}$ ontology language. This leads to the requirement that $\mathbf{L}_E[i] \leq \max(\mathbf{L}_C[i], \mathbf{L}_D[i])$ and $\min(\mathbf{U}_C[i], \mathbf{U}_D[i]) \leq \mathbf{U}_E[i]$ must hold for each embedding dimension i, where $\mathbf{L}_X, \mathbf{U}_X$ indicates the lower and upper bounds of the box associated to a concept X. Because the set $\mathcal{S} := \{(a,b,c) \in \mathbb{R}^3 \mid \min(a,b) \leq c\}$ is not convex[8], the aforementioned restrictions are not convex in general. Furthermore, recalling that optimal sets of convex functions are convex, it is not possible to devise a convex $f : \mathbb{R}^3 \to \mathbb{R}$ such that "$(a,b,c) \in \mathcal{S} \Leftrightarrow (a,b,c)$ is optimal for f" holds. In particular, absent extenuating circumstances, if the bounds of C, D, E are parameters to be optimized during learning, it is impossible to construct a nonnegative *convex* loss function that is zero if and only if "$C \sqcap D \sqsubseteq E$" is satisfied. Regarding *cone semantics*, although convex optimization is mentioned as one motivation for using cones in (Lütfü Özçep, Leemhuis, and Wolter, 2020), the authors do not explain how exactly convex optimization fits in the picture of their approach. In another work, the authors show how to use axis-aligned cones and pairs of unions of convex cones to solve certain multi-label classification problems (Leemhuis, Özçep, and Wolter, 2022) via SVMs. The approach described in (Leemhuis, Özçep, and Wolter, 2022) seems to be significantly different from the loss function minimization approach described in other papers. Furthermore, the Propositional $\mathcal{ALC}$ language used in (Leemhuis, Özçep, and Wolter, 2022) is different from DL-Lite$^{\mathcal{H}}$, which we consider in this work, since the latter features roles and inverses.

8 Conclusion and Future Work

We propose BoxLitE, a KB embedding method that allows for convex optimization and ensures the satisfaction of TBox axioms. We prove that for any DL-Lite$^{\mathcal{H}}$ KB, there is a faithful embedding solution that is a KB model. We implement and evaluate BoxLitE's convex problem formulation for KB embeddings in CVXPY and MOSEK. The results reveal that MOSEK finds a solution that is KB faithful and that satisfies the concept inclusions in the TBox for a prototypical ontology. Within a few seconds of solving time, MOSEK finds embedding solutions on subsets (F_v1 to F_v4) of a real-world KB that achieve good link prediction results. In the future, we would like to consider more efficient methods for the HPO and evaluation. Also, we want to study how to ensure other theoretical properties in convex KBE methods. As languages that allow for conjunctions on the left of inclusions often lead to nonconvexity, another line lies in studying sound nonconvex KBE approaches that can ensure the construction of faithful models. Finally, we would like to investigate, using nonconvex tools, the effect of pushing negative samples outside the concept box into an 'unknown' truth state, and how this affects prediction results.

Acknowledgements

This work was supported by the "Strategic Research Projects" grant from ROIS (Research Organization of Information and Systems). Bruno F. Lourenço's work was partially supported by the JSPS Grant-in-Aid for Early-Career Scientists 23K16844. Ana Ozaki was supported by the Research Council of Norway, projects (316022, 322480) and Integreat - Norwegian Centre for knowledge-driven machine learning (332645). Emanuel Sallinger's and Hesham Morgan's work on this paper was funded by the Vienna Science and Technology Fund (WWTF) [Grant ID: 10.47379/VRG18013, 10.47379/ICT25032, 10.47379/NXT22018, 10.47379/ICT2201, 10.47379/DCDH001], and by the Austrian Science Fund (FWF) 10.55776/COE12.

[8]It suffices to observe that $(1,0,0),(0,1,0) \in \mathcal{S}$, but $0.5(1,0,0) + 0.5(0,1,0) = (0.5,0.5,0) \notin \mathcal{S}$.

AI Declaration

Gen AI tools were only used to help find grammatical mistakes and to aid in the process of debugging the code.

References

Abboud, R.; Ceylan, İ. İ.; Lukasiewicz, T.; and Salvatori, T. 2020. BoxE: A box embedding model for knowledge base completion. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *NeurIPS*.

Agrawal, A.; Verschueren, R.; Diamond, S.; and Boyd, S. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5(1):42–60.

Ali, M.; Berrendorf, M.; Hoyt, C. T.; Vermue, L.; Sharifzadeh, S.; Tresp, V.; and Lehmann, J. 2021. PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings. *Journal of Machine Learning Research* 22(82):1–6.

ApS, M. 2026. *MOSEK Optimizer API for Python 11.1.3*.

Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyaschev, M. 2009. The DL-Lite family and relations. *J. Artif. Intell. Res.* 36:1–69.

Balazevic, I.; Allen, C.; and Hospedales, T. 2019. TuckER: Tensor factorization for knowledge graph completion. In *EMNLP-IJCNLP*, 5185–5194. Association for Computational Linguistics.

Beck, A. 2017. *First-Order Methods in Optimization*. SIAM.

Ben-Tal, A., and Nemirovski, A. 2001. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. SIAM.

Bolte, J., and Pauwels, E. 2020. Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning. *Mathematical Programming* 188(1):19–51.

Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In Burges, C. J. C.; Bottou, L.; Ghahramani, Z.; and Weinberger, K. Q., eds., *27th Annual Conference on Neural Information Processing Systems*, 2787–2795.

Bourgaux, C.; Guimarães, R.; Koudijs, R.; Lacerda, V.; and Ozaki, A. 2024. Knowledge base embeddings: Semantics and theoretical properties. In Marquis, P.; Ortiz, M.; and Pagnucco, M., eds., *KR*.

Bourgaux, C.; Ozaki, A.; and Pan, J. Z. 2021. Geometric models for (temporally) attributed description logics. In Homola, M.; Ryzhikov, V.; and Schmidt, R. A., eds., *(DL,* volume 2954 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge, England: Cambridge University Press.

Davis, D.; Drusvyatskiy, D.; Kakade, S.; and Lee, J. D. 2019. Stochastic subgradient method converges on tame functions. *Foundations of Computational Mathematics* 20(1):119–154.

Delfour, M. C., and Zolésio, J. P. 2011. *Shapes and Geometries: Metrics, Analysis, Differential Calculus, and Optimization, Second Edition*. Society for Industrial and Applied Mathematics.

Diamond, S., and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17(83):1–5.

Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.

Gutiérrez-Basulto, V., and Schockaert, S. 2018. From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules. In Thielscher, M.; Toni, F.; and Wolter, F., eds., *KR*, 379–388. AAAI Press.

Hiriart-Urruty, J.-B., and Lemaréchal, C. 1993a. *Convex Analysis and Minimization Algorithms I: Fundamentals*. Springer Berlin Heidelberg.

Hiriart-Urruty, J.-B., and Lemaréchal, C. 1993b. *Convex Analysis and Minimization Algorithms II: Advanced Theory and Bundle Methods*. Berlin Heidelberg: Springer.

Imenes, A.; Guimarães, R.; and Ozaki, A. 2023. Marrying query rewriting and knowledge graph embeddings. In *RuleML+RR*, 126–140. Berlin, Heidelberg: Springer-Verlag.

Jackermeier, M.; Chen, J.; and Horrocks, I. 2024. Dual box embeddings for the description logic EL++. In *Proceedings of the ACM Web Conference*, WWW, 2250–2258. Association for Computing Machinery.

Kazemi, S. M., and Poole, D. 2018. Simple embedding for link prediction in knowledge graphs. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *NeurIPS*, 4289–4300.

Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In Bengio, Y., and LeCun, Y., eds., *ICLR*.

Kontchakov, R.; Lutz, C.; Toman, D.; Wolter, F.; and Zakharyaschev, M. 2010. The combined approach to query answering in dl-lite. In Lin, F.; Sattler, U.; and Truszczynski, M., eds., *KR*. AAAI Press.

Kulmanov, M.; Liu-Wei, W.; Yan, Y.; and Hoehndorf, R. 2019. EL embeddings: Geometric construction of models for the description logic EL++. In Kraus, S., ed., *IJCAI*, 6103–6109. ijcai.org.

Lacerda, V.; Ozaki, A.; and Guimarães, R. 2024a. Faithel: Strongly tbox faithful knowledge base embeddings for *E*. In Kirrane, S.; Simkus, M.; Soylu, A.; and Roman, D., eds., *RuleML+RR*, volume 15183 of *Lecture Notes in Computer Science*, 191–199. Springer.

Lacerda, V.; Ozaki, A.; and Guimarães, R. 2024b. Strong faithfulness for ELH ontology embeddings. *TGDK* 2(3):2:1–2:29.

Leemhuis, M.; Özçep, O. L.; and Wolter, D. 2022. Learning with cone-based geometric models and orthologics. *Annals of Mathematics and Artificial Intelligence* 90(11–12):1159–1195.

Leskovec, J.; Kleinberg, J. M.; and Faloutsos, C. 2005. Graphs over time: densification laws, shrinking diameters and possible explanations. In Grossman, R.; Bayardo, R. J.; and Bennett, K. P., eds., *ACM SIGKDD*, 177–187. ACM.

Li, W.; Zheng, X.; Gao, H.; Ji, Q.; and Qi, G. 2022. Cosine-based embedding for completing lightweight schematic knowledge in dl-litecore. *Applied Sciences* 12(20).

Lobo, M. S.; Vandenberghe, L.; Boyd, S.; and Lebret, H. 1998. Applications of second-order cone programming. *Linear Algebra and its Applications* 284(1–3):193–228.

Lu, H., and Hu, H. 2020. Dense: An enhanced non-abelian group representation for knowledge graph embedding. *CoRR* abs/2008.04548.

Luo, H.; Wang, X.; and Lukens, B. 2018. Variational analysis on the signed distance functions. *Journal of Optimization Theory and Applications* 180(3):751–774.

Lütfü Özçep, O.; Leemhuis, M.; and Wolter, D. 2020. Cone semantics for logics with negation. In Bessiere, C., ed., *IJCAI*, 1820–1826. International Joint Conferences on Artificial Intelligence Organization. Main track.

Mittelmann, H. 2026. Large second order cone benchmark. https://plato.asu.edu/ftp/socp.html. [Online; accessed 17-February-2026].

Mondal, S.; Bhatia, S.; and Mutharaju, R. 2021. Emel++: Embeddings for EL++ description logic. In Martin, A.; Hinkelmann, K.; Fill, H.; Gerber, A.; Lenat, D.; Stolle, R.; and van Harmelen, F., eds., *AAAI-MAKE*, volume 2846 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Morgan, H.; Pavlović, A.; Ozaki, A.; Lourenço, B. F.; and Sallinger, E. 2026. Convexity for trustworthy knowledge graph embeddings: From interpretability to reliability. In *Proceedings of the 19th Forschungsforum der österreichischen Fachhochschulen (FHK 2026), Graz, Austria, April 15 - April 16, 2026*.

Nesterov, Y., and Nemirovskii, A. 1994. *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM.

Pavlović, A., and Sallinger, E. 2023a. Building bridges: Knowledge graph embeddings respecting logical rules. In Kimelfeld, B.; Martinez, M. V.; and Angles, R., eds., *Proceedings of the 15th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2023), Santiago de Chile, Chile, May 22-26, 2023*, volume 3409 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Pavlović, A., and Sallinger, E. 2023b. ExpressivE: A spatio-functional embedding for knowledge graph completion. In *the 11th International Conference on Learning Representations (ICLR 2023)*. OpenReview.net.

Pavlović, A., and Sallinger, E. 2024a. Raising the efficiency of knowledge graph embeddings while respecting logical rules. In Montoya, G.; Sallinger, E.; and Vargas-Solar, G., eds., *Proceedings of the 16th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2024), Mexico City, Mexico, September 30 - October 4, 2024*, CEUR Workshop Proceedings. CEUR-WS.org.

Pavlović, A., and Sallinger, E. 2024b. SpeedE: Euclidean geometric knowledge graph embedding strikes back. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Findings of the Association for Computational Linguistics: NAACL 2024*, 69–92. Mexico City, Mexico: Association for Computational Linguistics.

Pavlović, A.; Sallinger, E.; and Schockaert, S. 2025. Faithful Differentiable Reasoning with Reshuffled Region-based Embeddings. In *Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning*, 489–499.

Peng, X.; Tang, Z.; Kulmanov, M.; Niu, K.; and Hoehndorf, R. 2022. Description logic EL++ embeddings with intersectional closure. *CoRR* abs/2202.14018.

PyTorch. 2026. Pytorch - autograd mechanics. https://docs.pytorch.org/docs/stable/notes/autograd.html#gradients-for-non-differentiable-functions. [Online; accessed 17-February-2026].

Rockafellar, R. T. 1997. *Convex Analysis* . Princeton University Press.

Sturm, J. F. 1999. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software* 11(1-4):625–653.

Sun, Z.; Deng, Z.; Nie, J.; and Tang, J. 2019. RotatE: Knowledge graph embedding by relational rotation in complex space. In *ICLR*. OpenReview.net.

Trouillon, T.; Welbl, J.; Riedel, S.; Gaussier, É.; and Bouchard, G. 2016. Complex embeddings for simple link prediction. In Balcan, M., and Weinberger, K. Q., eds., *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, 2071–2080. JMLR.org.

Tütüncü, R. H.; Toh, K. C.; and Todd, M. J. 2003. Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming* 95(2):189–217.

Xiong, B.; Potyka, N.; Tran, T.; Nayyeri, M.; and Staab, S. 2022. Faithful embeddings for E^{++} knowledge bases. In Sattler, U.; Hogan, A.; Keet, C. M.; Presutti, V.; Almeida, J. P. A.; Takeda, H.; Monnin, P.; Pirrò, G.; and d'Amato, C., eds., *ISWC*, volume 13489 of *Lecture Notes in Computer Science*, 22–38. Springer.

Yang, B.; Yih, W.; He, X.; Gao, J.; and Deng, L. 2015. Embedding entities and relations for learning and inference in knowledge bases. In Bengio, Y., and LeCun, Y., eds., *ICLR*.

Yang, H.; Chen, J.; and Sattler, U. 2025. Transbox: El++-closed ontology embedding. In *Proceedings of the ACM on Web Conference*, WWW, 22–34. Association for Computing Machinery.

A Supplemental Material

This supplemental material contains additional information on the experimental setup, theoretical and empirical results, and complete proofs for each corollary, proposition, and theorem. Specifically, Section B introduces the semantics of DL-Lite$^{\mathcal{H}}$ for the convenience of the reader. Next, Section C provides all proofs for the theoretical results in Section 2. Afterwards, Section D provides proofs for the theoretical results of Section 4. Moreover, Section E formulates in detail BoxLitE's convex optimization problem and provides the proofs for each theoretical result of Section 5. Then, Section F provides additional information on the sizes of the optimization problem in the experiments. Section G provides additional information on the experimental setup, including implementation details and a discussion on the creation of the datasets (F_v1-4), the training setup, hyperparameter optimization, evaluation protocol, and used metrics.

B Basic Definitions: DL-Lite$^{\mathcal{H}}$ Semantics

For the convenience of the reader, here we provide the definition of the semantics for DL-Lite$^{\mathcal{H}}$, which is standard and can be found in references such as (Artale et al., 2009).

An *interpretation* $\mathcal{I}$ is a pair $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set, called the *domain* of $\mathcal{I}$, and $\cdot^{\mathcal{I}}$ is a function that assigns a subset $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ of the domain to each $A \in \mathsf{N_C}$, a binary relation $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ over the domain to each $R \in \mathsf{N_R}$, and an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ to each $a \in \mathsf{N_I}$. We extend $\cdot^{\mathcal{I}}$ to role and concept expressions as follows:

$$(\neg B)^{\mathcal{I}} := \Delta^{\mathcal{I}} \setminus B^{\mathcal{I}};$$
$$(R^{-})^{\mathcal{I}} := \{(e,d) \mid (d,e) \in R^{\mathcal{I}}\};$$
$$(\exists S)^{\mathcal{I}} := \{d \mid \exists e \in \Delta^{\mathcal{I}} \text{ such that } (d,e) \in S^{\mathcal{I}}\}.$$

We say that an interpretation $\mathcal{I}$ *satisfies*

- a role inclusion $S \sqsubseteq T$ iff $S^{\mathcal{I}} \subseteq T^{\mathcal{I}}$;
- a concept inclusion $B \sqsubseteq C$ iff $B^{\mathcal{I}} \subseteq C^{\mathcal{I}}$;
- a role assertion $R(a,b)$ iff $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$; and
- a concept assertion $D(a)$ iff $a^{\mathcal{I}} \in D^{\mathcal{I}}$.

In this work $\mathsf{N_C}$, $\mathsf{N_R}$, and $\mathsf{N_I}$ are all *finite sets*, considered to be the relevant symbols to express KBs. While in the Description Logic literature these sets are often assumed to be countably infinite, in the KB embedding literature, they are usually assumed to be finite, e.g., (Xiong et al., 2022) and that is what we adopt here. In what follows, we denote the size of a finite set V by $|V|$.

C Proofs for Section 2

Here, we provide proofs for the results in Section 2. The canonical interpretations found in the literature, e.g., (Kontchakov et al., 2010) are usually designed for query answering and may not satisfy the (concept/role) inclusions that are entailed by the TBox or falsify those inclusions that are not entailed. Since satisfying the TBox is important in our work for establishing faithfulness results later, we provided our own definition of the canonical model (Definition 1) and now provide the full proof of Theorem 1.

Theorem 1. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB and let $\mathcal{I}_{\mathcal{K}}$ be the canonical model of $\mathcal{K}$. Then, for all DL-Lite$^{\mathcal{H}}$ axioms α, we have that $\mathcal{I}_{\mathcal{K}} \models \alpha$ iff $\mathcal{K} \models \alpha$.*

Proof. In the following, assume $A, B \in \mathsf{N_C}$ and $R, S \in \mathsf{N_R}$.

Claim 1. $\mathcal{I}_{\mathcal{K}} \models A \sqsubseteq B$ *iff* $\mathcal{K} \models A \sqsubseteq B$.

Proof. **Assume** $\mathcal{K} \models A \sqsubseteq B$**.** We make a case distinction based on the elements in $\Delta^{\mathcal{I}_{\mathcal{K}}} := \mathsf{N_I} \cup \Delta_{\mathcal{K}}$.

- $a \in \mathsf{N_I}$: Assume $a \in A^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $a \in A^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models A(a)$. By assumption, $\mathcal{K} \models A \sqsubseteq B$, so $\mathcal{K} \models B(a)$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $a \in B^{\mathcal{I}_{\mathcal{K}}}$. Since a was an arbitrary element of $\mathsf{N_I}$ this holds for all elements of this kind.
- $c_D \in \Delta_{\mathcal{K}}$: Assume $c_D \in A^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $c_D \in A^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models D \sqsubseteq A$. By assumption, $\mathcal{K} \models A \sqsubseteq B$, so $\mathcal{K} \models D \sqsubseteq B$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $c_D \in B^{\mathcal{I}_{\mathcal{K}}}$. Since c_D was an arbitrary element of $\Delta_{\mathcal{K}}$ this holds for all elements of this kind.

We have thus shown that $\mathcal{I}_{\mathcal{K}} \models A \sqsubseteq B$.

Now, assume $\mathcal{K} \not\models A \sqsubseteq B$**.** We show that $\mathcal{I}_{\mathcal{K}} \not\models A \sqsubseteq B$. If $\mathcal{K} \not\models A \sqsubseteq B$ then there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}$ with $A^{\mathcal{I}}$ non-empty. This means that A is satisfiable w.r.t. $\mathcal{K}$ and thus $c_A \in \Delta_{\mathcal{K}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $c_A \in A^{\mathcal{I}_{\mathcal{K}}}$ since $\mathcal{K} \models A \sqsubseteq A$ holds trivially. We now argue that $c_A \notin B^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, an element of the form c_D is in $B^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models D \sqsubseteq B$. By assumption $\mathcal{K} \not\models A \sqsubseteq B$. So c_A is not in $B^{\mathcal{I}_{\mathcal{K}}}$. ☐

Claim 2. $\mathcal{I}_{\mathcal{K}} \models A \sqsubseteq \neg B$ *iff* $\mathcal{K} \models A \sqsubseteq \neg B$.

Proof. **Assume** $\mathcal{K} \models A \sqsubseteq \neg B$**.** We make a case distinction based on the elements in $\Delta^{\mathcal{I}_{\mathcal{K}}} := \mathsf{N_I} \cup \Delta_{\mathcal{K}}$.

- $a \in \mathsf{N_I}$: Assume $a \in A^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $a \in A^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models A(a)$. By assumption, $\mathcal{K} \models A \sqsubseteq \neg B$. Also, by assumption $\mathcal{K}$ is satisfiable, meaning that $\mathcal{K} \not\models B(a)$. Then, by definition of $\mathcal{I}_{\mathcal{K}}$, $a \notin B^{\mathcal{I}_{\mathcal{K}}}$, that is, $a \in (\neg B)^{\mathcal{I}_{\mathcal{K}}}$. As a was an arbitrary element of $\mathsf{N_I}$ this holds for all elements of this kind.
- $c_D \in \Delta_{\mathcal{K}}$: Assume $c_D \in A^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $c_D \in A^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models D \sqsubseteq A$. By assumption, $\mathcal{K} \models A \sqsubseteq \neg B$, so $\mathcal{K} \models D \sqsubseteq \neg B$. As $c_D \in \Delta_{\mathcal{K}}$, by definition of $\Delta_{\mathcal{K}}$, D is satisfiable w.r.t $\mathcal{K}$. So $\mathcal{K} \not\models D \sqsubseteq B$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $c_D \in (\neg B)^{\mathcal{I}_{\mathcal{K}}}$. Since c_D was an arbitrary element of $\Delta_{\mathcal{K}}$ this holds for all elements of this kind.

We have thus shown that $\mathcal{I}_{\mathcal{K}} \models A \sqsubseteq \neg B$.

Now, assume $\mathcal{K} \not\models A \sqsubseteq \neg B$**.** We show that $\mathcal{I}_{\mathcal{K}} \not\models A \sqsubseteq \neg B$. If $\mathcal{K} \not\models A \sqsubseteq \neg B$ then there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}$ with $A^{\mathcal{I}} \cap B^{\mathcal{I}}$ non-empty. This means that $A \sqcap B$ is satisfiable w.r.t. $\mathcal{K}$ and thus $c_{A \sqcap B} \in \Delta_{\mathcal{K}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $c_{A \sqcap B} \in A^{\mathcal{I}_{\mathcal{K}}}$ and $c_{A \sqcap B} \in B^{\mathcal{I}_{\mathcal{K}}}$ since $\mathcal{K} \models A \sqcap B \sqsubseteq A$ and $\mathcal{K} \models A \sqcap B \sqsubseteq B$. So $\mathcal{I}_{\mathcal{K}} \not\models A \sqsubseteq \neg B$. ☐

Claim 3. $\mathcal{I}_{\mathcal{K}} \models R \sqsubseteq S$ *iff* $\mathcal{K} \models R \sqsubseteq S$.

Proof. **Assume** $\mathcal{K} \models R \sqsubseteq S$**.** We make a case distinction based on the elements in $\Delta^{\mathcal{I}_{\mathcal{K}}}$ and how they can be related in the extension of a role name in the definition of $\mathcal{I}_{\mathcal{K}}$.

- $(a, b) \in \mathsf{N_I} \times \mathsf{N_I}$: Assume $(a, b) \in R^{\mathcal{I}_{\mathcal{K}}}$. We first argue that in this case $\mathcal{K} \models R(a, b)$. By definition of $\mathcal{I}_{\mathcal{K}}$, $(a, b) \in R^{\mathcal{I}_{\mathcal{K}}}$ iff $\mathcal{K} \models R(a, b)$. Since by assumption $\mathcal{K} \models R \sqsubseteq S$ we have that $\mathcal{K} \models S(a, b)$, so $(a, b) \in S^{\mathcal{I}_{\mathcal{K}}}$. Since (a, b) was an arbitrary pair in $\mathsf{N_I} \times \mathsf{N_I}$, this holds for all such kinds of pairs.
- $(a, c_{\exists R'}) \in \mathsf{N_I} \times \Delta_{\mathcal{K}}$: Assume $(a, c_{\exists R'}) \in R^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $\mathcal{K} \models \exists \overline{R'}(a)$ and $\mathcal{K} \models \overline{R'} \sqsubseteq R$. By assumption $\mathcal{K} \models R \sqsubseteq S$. So $\mathcal{K} \models \overline{R'} \sqsubseteq S$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, we have that $(a, c_{\exists R'}) \in S^{\mathcal{I}_{\mathcal{K}}}$. Since $(a, c_{\exists R'})$ was an arbitrary pair of this format in $\mathsf{N_I} \times \Delta_{\mathcal{K}}$, this holds for all such kinds of pairs.
- $(c_{\exists R'}, a) \in \Delta_{\mathcal{K}} \times \mathsf{N_I}$: Assume $(c_{\exists R'}, a) \in R^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, $\mathcal{K} \models \exists \overline{R'}(a)$ and $\mathcal{K} \models R' \sqsubseteq R$. By assumption $\mathcal{K} \models R \sqsubseteq S$. So $\mathcal{K} \models R' \sqsubseteq S$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, we have that $(c_{\exists R'}, a) \in S^{\mathcal{I}_{\mathcal{K}}}$. Since $(c_{\exists R'}, a)$ was an arbitrary pair of this format in $\mathsf{N_I} \times \Delta_{\mathcal{K}}$, this argument can be applied for all such kinds of pairs.
- $(c_{\exists R'}, c_{\exists \overline{R'}}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}}$: Assume $(c_{\exists R'}, c_{\exists \overline{R'}}) \in R^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $\mathcal{K} \models R' \sqsubseteq R$. By assumption, $\mathcal{K} \models R \sqsubseteq S$, so $\mathcal{K} \models R' \sqsubseteq S$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $(c_{\exists R'}, c_{\exists \overline{R'}}) \in S^{\mathcal{I}_{\mathcal{K}}}$.
- $(c_D, c_{\exists R'}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}}$: Assume $(c_D, c_{\exists R'}) \in R^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $\mathcal{K} \models D \sqsubseteq \exists \overline{R'}$ and $\mathcal{K} \models \overline{R'} \sqsubseteq R$. By assumption, $\mathcal{K} \models R \sqsubseteq S$, so $\mathcal{K} \models \overline{R'} \sqsubseteq S$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $(c_D, c_{\exists R'}) \in S^{\mathcal{I}_{\mathcal{K}}}$.
- $(c_{\exists R'}, c_D) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}}$: Assume $(c_{\exists R'}, c_D) \in R^{\mathcal{I}_{\mathcal{K}}}$. By definition of $\mathcal{I}_{\mathcal{K}}$, we have that $\mathcal{K} \models D \sqsubseteq \exists \overline{R'}$ and $\mathcal{K} \models R' \sqsubseteq R$. By assumption, $\mathcal{K} \models R \sqsubseteq S$, so $\mathcal{K} \models R' \sqsubseteq S$. Then, again by definition of $\mathcal{I}_{\mathcal{K}}$, $(c_{\exists R'}, c_D) \in S^{\mathcal{I}_{\mathcal{K}}}$.

We have thus shown that $\mathcal{I}_{\mathcal{K}} \models R \sqsubseteq S$.

Now, assume $\mathcal{K} \not\models R \sqsubseteq S$**.** We show that $\mathcal{I}_\mathcal{K} \not\models R \sqsubseteq S$. By definition of $\mathcal{I}_\mathcal{K}$, we have that $\{(c_{\exists S}, c_{\exists \overline{S}}) \in \Delta_\mathcal{K} \times \Delta_\mathcal{K} \mid \mathcal{K} \models S \sqsubseteq R\} \subseteq R^{\mathcal{I}_\mathcal{K}}$. By taking $S = R$ (and since trivially $\mathcal{K} \models R \sqsubseteq R$), we have in particular that $(c_{\exists R}, c_{\exists \overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$. We now argue that $(c_{\exists R}, c_{\exists \overline{R}}) \notin S^{\mathcal{I}_\mathcal{K}}$. By definition of $\mathcal{I}_\mathcal{K}$, a pair of the form $(c_{\exists S'}, c_{\exists \overline{S'}})$ is in $S^{\mathcal{I}_\mathcal{K}}$ iff $\mathcal{K} \models S' \sqsubseteq S$. By assumption $\mathcal{K} \not\models R \sqsubseteq S$. So $(c_{\exists R}, c_{\exists \overline{R}}) \notin S^{\mathcal{I}_\mathcal{K}}$. □

Claim 4. *$\mathcal{I}_\mathcal{K} \models \exists R \sqsubseteq A$ iff $\mathcal{K} \models \exists R \sqsubseteq A$.*

Proof. **Assume** $\mathcal{K} \models \exists R \sqsubseteq A$**.** We make a case distinction based on the elements in $\Delta^{\mathcal{I}_\mathcal{K}} := \mathsf{N_I} \cup \Delta_\mathcal{K}$.

- $a \in \mathsf{N_I}$: Assume $a \in (\exists R)^{\mathcal{I}_\mathcal{K}}$. In this case, by definition of $\mathcal{I}_\mathcal{K}$, either (1) there is $b \in \mathsf{N_I}$ such that $(a, b) \in R^{\mathcal{I}_\mathcal{K}}$ or (2) there is $c_{\exists \overline{R}} \in \Delta_\mathcal{K}$ such that $(a, c_{\exists \overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$. In case (1), by definition of $\mathcal{I}_\mathcal{K}$, $(a, b) \in R^{\mathcal{I}_\mathcal{K}}$ implies that $\mathcal{K} \models R(a, b)$. Together with the assumption that $\mathcal{K} \models \exists R \sqsubseteq A$, this means that $\mathcal{K} \models A(a)$. Again by definition of $\mathcal{I}_\mathcal{K}$, we have that $a \in A^{\mathcal{I}_\mathcal{K}}$. In case (2), by definition of $\mathcal{I}_\mathcal{K}$, $(a, c_{\exists \overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$ implies that $\mathcal{K} \models \exists R(a)$. By assumption $\mathcal{K} \models \exists R \sqsubseteq A$, which means that $\mathcal{K} \models A(a)$. Again by definition of $\mathcal{I}_\mathcal{K}$, we have that $a \in A^{\mathcal{I}_\mathcal{K}}$. Since a was an arbitrary element in $\mathsf{N_I}$, this argument can be applied for all elements of this kind.
- $c_D \in \Delta_\mathcal{K}$: Assume $c_D \in (\exists R)^{\mathcal{I}_\mathcal{K}}$. We first show that $\mathcal{K} \models D \sqsubseteq \exists R$. If $c_D \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ then, by definition of $\mathcal{I}_\mathcal{K}$, either (1) there is $a \in \mathsf{N_I}$ such that $(c_D, a) \in R^{\mathcal{I}_\mathcal{K}}$ or (2) there is $c_{D'} \in \Delta_\mathcal{K}$ such that $(c_D, c_{D'}) \in R^{\mathcal{I}_\mathcal{K}}$. In case (1), by definition of $\mathcal{I}_\mathcal{K}$, $(c_D, a) \in R^{\mathcal{I}_\mathcal{K}}$ means that D is of the form $\exists S$, $\mathcal{K} \models S \sqsubseteq R$, and $\mathcal{K} \models \exists \overline{S}(a)$. So $\mathcal{K} \models D \sqsubseteq \exists R$. In case (2), by definition of $\mathcal{I}_\mathcal{K}$, there are three possibilities:

  (a) D is of the form $\exists S$ and D' is of the form $\exists \overline{S}$;

  (b) D' is of the form $\exists S$ and $\mathcal{K} \models D \sqsubseteq \exists \overline{S}$;

  (c) D is of the form $\exists S$ and $\mathcal{K} \models D' \sqsubseteq \exists \overline{S}$.

  In the sub-cases (a) and (c) we also have that $\mathcal{K} \models S \sqsubseteq R$. So $\mathcal{K} \models D \sqsubseteq \exists R$. In the sub-case (b), we have that $\mathcal{K} \models \overline{S} \sqsubseteq R$. Then again $\mathcal{K} \models D \sqsubseteq \exists R$. By assumption $\mathcal{K} \models \exists R \sqsubseteq A$, which means that $\mathcal{K} \models D \sqsubseteq A$. Then, $c_D \in A^{\mathcal{I}_\mathcal{K}}$, by definition of $\mathcal{I}_\mathcal{K}$. Since c_D was an arbitrary element in $\Delta_\mathcal{K}$, this argument can be applied for all elements of this kind.

We have thus shown that, for all elements d in $\Delta^{\mathcal{I}_\mathcal{K}}$, if $d \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ then $d \in A^{\mathcal{I}_\mathcal{K}}$. So $\mathcal{I}_\mathcal{K} \models \exists R \sqsubseteq A$.

Now, assume $\mathcal{K} \not\models \exists R \sqsubseteq A$**.** If $\mathcal{K} \not\models \exists R \sqsubseteq A$ then there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}$ with $(\exists R)^\mathcal{I}$ non-empty. This means that $\exists R$ is satisfiable w.r.t. $\mathcal{K}$ and thus $c_{\exists R} \in \Delta_\mathcal{K}$. We show that $\mathcal{I}_\mathcal{K} \not\models \exists R \sqsubseteq A$ by showing that $c_{\exists R} \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ but $c_{\exists R} \notin A^{\mathcal{I}_\mathcal{K}}$. By the definition of $\mathcal{I}_\mathcal{K}$, $(c_{\exists S}, c_{\exists \overline{S}}) \in R^{\mathcal{I}_\mathcal{K}}$ if $\mathcal{K} \models S \sqsubseteq R$, which is trivially the case for $S = R$. So $c_{\exists R} \in (\exists R)^{\mathcal{I}_\mathcal{K}}$. We now argue that $c_{\exists R} \notin A^{\mathcal{I}_\mathcal{K}}$. By definition of $\mathcal{I}_\mathcal{K}$, an element of the form c_D is in $A^{\mathcal{I}_\mathcal{K}}$ iff $\mathcal{K} \models D \sqsubseteq A$. By assumption $\mathcal{K} \not\models \exists R \sqsubseteq A$. So $c_{\exists R}$ is not in $A^{\mathcal{I}_\mathcal{K}}$. □

Claim 5. *$\mathcal{I}_\mathcal{K} \models A \sqsubseteq \exists R$ iff $\mathcal{K} \models A \sqsubseteq \exists R$.*

Proof. **Assume** $\mathcal{K} \models A \sqsubseteq \exists R$. We make a case distinction based on the elements in $\Delta^{\mathcal{I}_\mathcal{K}} := \mathsf{N_I} \cup \Delta_\mathcal{K}$.

- $a \in \mathsf{N_I}$: Assume $a \in A^{\mathcal{I}_\mathcal{K}}$. In this case, by definition of $\mathcal{I}_\mathcal{K}$, we have that $\mathcal{K} \models A(a)$. By assumption, $\mathcal{K} \models A \sqsubseteq \exists R$. So $\mathcal{K} \models \exists R(a)$. Then, by definition of $\mathcal{I}_\mathcal{K}$, $(a, c_{\exists \overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$. Thus, $a \in (\exists R)^{\mathcal{I}_\mathcal{K}}$, as required. Since a was an arbitrary element in $\mathsf{N_I}$, this argument can be applied for all elements of this kind.
- $c_D \in \Delta_\mathcal{K}$: Assume $c_D \in A^{\mathcal{I}_\mathcal{K}}$. In this case, by definition of $\mathcal{I}_\mathcal{K}$, we have that $\mathcal{K} \models D \sqsubseteq A$. By assumption, $\mathcal{K} \models A \sqsubseteq \exists R$, so $\mathcal{K} \models D \sqsubseteq \exists R$. Then, by definition of $\mathcal{I}_\mathcal{K}$, we have that $(c_D, c_{\exists \overline{S}}) \in R^{\mathcal{I}_\mathcal{K}}$ for $S \in \mathsf{N_R^-}$ such that $\mathcal{K} \models S \sqsubseteq R$. Taking $S = R$ this trivially holds, so $(c_D, c_{\exists \overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$. This means that $c_D \in (\exists R)^{\mathcal{I}_\mathcal{K}}$, as required. Since c_D was an arbitrary element in $\Delta_\mathcal{K}$, this argument can be applied for all elements of this kind.

We have thus shown that, for all elements d in $\Delta^{\mathcal{I}_\mathcal{K}}$, if $d \in A^{\mathcal{I}_\mathcal{K}}$ then $d \in (\exists R)^{\mathcal{I}_\mathcal{K}}$. So $\mathcal{I}_\mathcal{K} \models A \sqsubseteq \exists R$.

Now, assume $\mathcal{K} \not\models A \sqsubseteq \exists R$**.** If $\mathcal{K} \not\models A \sqsubseteq \exists R$ then there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}$ with $A^\mathcal{I}$ non-empty. This means that A is satisfiable w.r.t. $\mathcal{K}$ and thus $c_A \in \Delta_\mathcal{K}$. We show that $\mathcal{I}_\mathcal{K} \not\models A \sqsubseteq \exists R$ by showing that $c_A \in A^{\mathcal{I}_\mathcal{K}}$ but $c_A \notin (\exists R)^{\mathcal{I}_\mathcal{K}}$. By definition of $\mathcal{I}_\mathcal{K}$, an element of the form c_D is in $A^{\mathcal{I}_\mathcal{K}}$ iff $\mathcal{K} \models D \sqsubseteq A$, which is trivially the case for $D = A$. So $c_A \in A^{\mathcal{I}_\mathcal{K}}$. We now argue that $c_A \notin (\exists R)^{\mathcal{I}_\mathcal{K}}$. By definition of $\mathcal{I}_\mathcal{K}$, $c_A \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ iff $\mathcal{K} \models A \sqsubseteq \exists S$ and $\mathcal{K} \models S \sqsubseteq R$. By assumption $\mathcal{K} \not\models A \sqsubseteq \exists R$. So either $\mathcal{K} \not\models A \sqsubseteq \exists S$ or $\mathcal{K} \not\models S \sqsubseteq R$. Thus, $c_A \notin (\exists R)^{\mathcal{I}_\mathcal{K}}$. □

Claim 6. *$\mathcal{I}_\mathcal{K} \models A \sqsubseteq \neg\exists R$ iff $\mathcal{K} \models A \sqsubseteq \neg\exists R$.*

Proof. **Assume** $\mathcal{K} \models A \sqsubseteq \neg\exists R$. We make a case distinction based on the elements in $\Delta^{\mathcal{I}_\mathcal{K}} := \mathsf{N_I} \cup \Delta_\mathcal{K}$.

- $a \in \mathsf{N_I}$: Assume $a \in A^{\mathcal{I}_\mathcal{K}}$. In this case, by definition of $\mathcal{I}_\mathcal{K}$, we have that $\mathcal{K} \models A(a)$. By assumption, $\mathcal{K} \models A \sqsubseteq \neg\exists R$. As $\mathcal{K}$ is satisfiable, $\mathcal{K} \not\models \exists R(a)$. By definition of $\mathcal{I}_\mathcal{K}$, to show that $\mathcal{I}_\mathcal{K} \not\models \exists R(a)$, we need to rule out the following two cases, where $R' \in \mathsf{N_R}$.
  - $(a, b) \in R'^{\mathcal{I}_\mathcal{K}}$ and $R' = R$. Since this would imply $\mathcal{K} \models R'(a, b)$ and $\mathcal{K} \not\models \exists R(a)$, this cannot happen.
  - $(a, c_{\exists S}) \in R'^{\mathcal{I}_\mathcal{K}}$ and $R' = R$. Since $\mathcal{K} \not\models \exists R(a)$, there is no role S such that $(a, c_{\exists S}) \in R'^{\mathcal{I}_\mathcal{K}}$, $\mathcal{K} \models \exists \overline{S}(a)$, $\mathcal{K} \models \overline{S} \sqsubseteq R$. So this case cannot happen because otherwise we would have that $\mathcal{K} \models \exists R(a)$.
  - $(c_{\exists S}, a) \in R'^{\mathcal{I}_\mathcal{K}}$ and $R' = R^-$. Since $\mathcal{K} \not\models \exists R(a)$, there is no role S such that $(c_{\exists S}, a) \in R'^{\mathcal{I}_\mathcal{K}}$, $\mathcal{K} \models \exists \overline{S}(a)$, and $\mathcal{K} \models S \sqsubseteq R$. So this case cannot happen because otherwise $\mathcal{K} \models \exists R(a)$ would hold.

  Then, by definition of $\mathcal{I}_\mathcal{K}$, we have that $\mathcal{I}_\mathcal{K} \not\models \exists R(a)$. Since a was an arbitrary element in $\mathsf{N_I}$, this argument can be applied for all elements of this kind.

- $c_D \in \Delta_\mathcal{K}$: Assume $c_D \in A^{\mathcal{I}_\mathcal{K}}$. In this case, by definition of $\mathcal{I}_\mathcal{K}$, we have that $\mathcal{K} \models D \sqsubseteq A$. By assumption, $\mathcal{K} \models A \sqsubseteq \neg\exists R$, so $\mathcal{K} \models D \sqsubseteq \neg\exists R$. In the second item of the argument of Claim 4, we have shown that if $c_D \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ then $\mathcal{K} \models D \sqsubseteq \exists R$. As $c_D \in \Delta_\mathcal{K}$, we have that D is satisfiable. So $\mathcal{K} \not\models D \sqsubseteq \exists R$. Then, by contrapositive, we have that $c_D \notin (\exists R)^{\mathcal{I}_\mathcal{K}}$.

We have thus shown that, for all elements d in $\Delta^{\mathcal{I}_\mathcal{K}}$, if $d \in A^{\mathcal{I}_\mathcal{K}}$ then $d \in (\neg\exists R)^{\mathcal{I}_\mathcal{K}}$. So $\mathcal{I}_\mathcal{K} \models A \sqsubseteq \neg\exists R$.

Now, assume $\mathcal{K} \not\models A \sqsubseteq \neg\exists R$. If $\mathcal{K} \not\models A \sqsubseteq \neg\exists R$ then there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}$ with $A^\mathcal{I} \cap (\exists R)^\mathcal{I}$ non-empty. This means that $A \sqcap \exists R$ is satisfiable w.r.t. $\mathcal{K}$ and thus $c_{A\sqcap\exists R} \in \Delta_\mathcal{K}$. By definition of $\mathcal{I}_\mathcal{K}$, we have that $c_{A\sqcap\exists R} \in A^{\mathcal{I}_\mathcal{K}}$ and $c_{A\sqcap\exists R} \in (\exists R)^{\mathcal{I}_\mathcal{K}}$ since $\mathcal{K} \models A \sqcap \exists R \sqsubseteq A$ and $\mathcal{K} \models A \sqcap \exists R \sqsubseteq \exists R$ (in more details, we have that $(c_{A\sqcap\exists R}, c_{\exists R^-}) \in R^{\mathcal{I}_\mathcal{K}}$). So $\mathcal{I}_\mathcal{K} \not\models A \sqsubseteq \neg\exists R$. □

Claim 7. *$\mathcal{I}_\mathcal{K} \models \exists R \sqsubseteq \neg A$ iff $\mathcal{K} \models \exists R \sqsubseteq \neg A$.*

Proof. Since $\mathcal{K} \models \exists R \sqsubseteq \neg A$ iff $\mathcal{K} \models A \sqsubseteq \neg\exists R$ and $\mathcal{I}_\mathcal{K} \models \exists R \sqsubseteq \neg A$ iff $\mathcal{I}_\mathcal{K} \models A \sqsubseteq \neg\exists R$, this claim follows from Claim 6. □

Regarding the assertions, we have that $\mathcal{I}_\mathcal{K} \models R(a,b)$ iff $\mathcal{K} \models R(a,b)$ directly follows from the definition of $\mathcal{I}_\mathcal{K}$ (Definition 1). The same holds for $\mathcal{I}_\mathcal{K} \models A(a)$ iff $\mathcal{K} \models A(a)$. It remains to argue about assertions of the form $\exists R'(a)$ (where R' is a role name or its inverse).

Claim 8. *$\mathcal{I}_\mathcal{K} \models \exists R'(a)$ iff $\mathcal{K} \models \exists R'(a)$.*

Proof. We first show that if $\mathcal{K} \models \exists R'(a)$ then $\mathcal{I}_\mathcal{K} \models \exists R'(a)$. We make a case distinction.

- $R' = R$. If $\mathcal{K} \models \exists R(a)$ and $R \in \mathsf{N_R}$ then $(a, c_{\exists\overline{R}}) \in R^{\mathcal{I}_\mathcal{K}}$ (take $\overline{S} = R$ in Definition 1). So $\mathcal{I}_\mathcal{K} \models \exists R(a)$.
- $R' = R^-$. If $\mathcal{K} \models \exists R^-(a)$ and $R \in \mathsf{N_R}$ then $(c_{\exists R}, a) \in R^{\mathcal{I}_\mathcal{K}}$ (take $\overline{S} = R^-$ in Definition 1). So $\mathcal{I}_\mathcal{K} \models \exists R^-(a)$.

Conversely, assume $\mathcal{I}_\mathcal{K} \models \exists R'(a)$. There are two cases.

- $R' = R$ and $(a, c_{\exists S}) \in R^{\mathcal{I}_\mathcal{K}}$ for some S such that $\mathcal{K} \models \exists\overline{S}(a)$ and $\mathcal{K} \models \overline{S} \sqsubseteq R$. In this case, $\mathcal{K} \models \exists R(a)$. Then $\mathcal{K} \models \exists R'(a)$ since here $R' = R$.
- $R' = R^-$ and $(c_{\exists S}, a) \in R^{\mathcal{I}_\mathcal{K}}$ for some S such that $\mathcal{K} \models \exists\overline{S}(a)$ and $\mathcal{K} \models S \sqsubseteq R$. In this case, note that $(c_{\exists S}, a) \in R^{\mathcal{I}_\mathcal{K}}$ means $\mathcal{I}_\mathcal{K} \models \exists R^-(a)$. As $\mathcal{K} \models \exists\overline{S}(a)$ and $\mathcal{K} \models S \sqsubseteq R$ we have that $\mathcal{K} \models \exists R^-(a)$ so $\mathcal{K} \models \exists R'(a)$ since here $R' = R^-$.

□

We have completed now our proof with all DL-Lite$^\mathcal{H}$ axioms. □

Corollary 3. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^\mathcal{H}$ KB and let $\mathcal{I}'_\mathcal{K}$ be the result of modifying the canonical model $\mathcal{I}_\mathcal{K}$ of $\mathcal{K}$ by setting $\Delta_\mathcal{K}$ as $\{c_{\exists R}, c_{\exists R^-} \mid R \in \mathsf{N_R}\}$. Then, for all DL-Lite$^\mathcal{H}$ axioms α, we have that if $\mathcal{K} \models \alpha$ then $\mathcal{I}_\mathcal{K} \models \alpha$.*

Proof. The proof is the same as in Theorem 1, except that since here we only want one direction of the theorem (showing that the interpretation satisfies $\mathcal{K}$) we do not require all the elements of $\Delta_\mathcal{K}$ used in Theorem 1. □

Theorem 2. *For any box interpretation η:*

i) for all $C \in \mathsf{N_C^\exists}$, $\overline{\eta(C)} \in \mathsf{Box}$;

ii) for all $C \in \mathsf{N_C^\exists}$, $\overline{\overline{\eta(C)}} = \eta(C)$;

iii) for all $C, D \in \mathsf{N_C^\exists}$, if $\eta(C) \subseteq \eta(D)$ then $\overline{\eta(D)} \subseteq \overline{\eta(C)}$.

Proof. (Condition i)) What is to be shown is that for all concept embeddings $\eta(C) \in \mathsf{Box}$ with $C \in \mathsf{N_C^\exists}$, there is a complement $\overline{\eta(C)} \in \mathsf{Box}$. By the definition of the complement of boxes, we know that $\overline{\eta(C)}$ is:

$$\{\mathbf{x} \mid (-\mathbf{s_\Omega} - \mathbf{L_C} + \boldsymbol{\epsilon}) \leq_d \mathbf{x} \leq_d (\mathbf{s_\Omega} - \mathbf{U_C} - \boldsymbol{\epsilon}), \mathbf{x} \in \mathbb{R}^d\}.$$

Given Definition 2 and Definition 3, the bounds of the complement box $\overline{\eta(C)}$ are $\mathbf{L_{\neg C}} = -\mathbf{s_\Omega} - \mathbf{L_C}$ and $\mathbf{U_{\neg C}} = \mathbf{s_\Omega} - \mathbf{U_C}$. The negation of a box $\overline{\eta(C)}$ is again a box, since: (i) we can compute the width of $\overline{\eta(C)}$ as $\mathbf{U_{\neg C}} - \mathbf{L_{\neg C}} = (\mathbf{s_\Omega} - \mathbf{U_C}) - (-\mathbf{s_\Omega} - \mathbf{L_C}) = 2\mathbf{s_\Omega} - (\mathbf{U_C} - \mathbf{L_C})$ and (ii) by the definition of the set Box of boxes we know that $\mathbf{0} \leq_d (\mathbf{U_C} - \mathbf{L_C}) \leq_d 2\mathbf{s_\Omega}$. From (i) and (ii) we have that $\mathbf{0} \leq_d 2\mathbf{s_\Omega} - (\mathbf{U_C} - \mathbf{L_C}) \leq_d 2\mathbf{s_\Omega}$, proving that $\overline{\eta(C)} \in \mathsf{Box}$. Thus, we have shown that for all concept embeddings $\eta(C) \in \mathsf{Box}$ with $C \in \mathsf{N_C^\exists}$, it holds that $\overline{\eta(C)} \in \mathsf{Box}$.

(Condition ii)) We show that $\overline{\overline{\eta(C)}} = \eta(C)$ for all $\eta(C) \in \mathsf{Box}$ with $C \in \mathsf{N_C^\exists}$. We have that $\overline{\overline{\eta(C)}}$ is

$$\begin{aligned}
&= \overline{\overline{\{\mathbf{x} \mid \mathbf{L_C} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U_C} - \boldsymbol{\epsilon}, \mathbf{x} \in \mathbb{R}^d\}}} \\
&= \overline{\{\mathbf{x} \mid (-\mathbf{s_\Omega} - \mathbf{L_C} + \boldsymbol{\epsilon}) \leq_d \mathbf{x} \leq_d (\mathbf{s_\Omega} - \mathbf{U_C} - \boldsymbol{\epsilon}), \mathbf{x} \in \mathbb{R}^d\}} \\
&= \{\mathbf{x} \mid (-\mathbf{s_\Omega} - (-\mathbf{s_\Omega} - \mathbf{L_C}) + \boldsymbol{\epsilon}) \leq_d \mathbf{x} \leq_d \\
&\quad (\mathbf{s_\Omega} - (\mathbf{s_\Omega} - \mathbf{U_C}) - \boldsymbol{\epsilon}), \mathbf{x} \in \mathbb{R}^d\} \\
&= \{\mathbf{x} \mid \mathbf{L_C} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U_C} - \boldsymbol{\epsilon}, \mathbf{x} \in \mathbb{R}^d\} \\
&= \eta(C).
\end{aligned}$$

(Condition iii)) We show that if $\eta(C) \subseteq \eta(D)$ then $\overline{\eta(D)} \subseteq \overline{\eta(C)}$, for all $\eta(C), \eta(D) \in \mathsf{Box}$ with $C, D \in \mathsf{N_C^\exists}$:

$$\begin{aligned}
&\eta(C) \subseteq \eta(D) \\
&\Leftrightarrow (\mathbf{L_D} \leq_d \mathbf{L_C}) \wedge (\mathbf{U_C} \leq_d \mathbf{U_D}) \\
&\Leftrightarrow ((\mathbf{s_\Omega} + \mathbf{L_D}) \leq_d (\mathbf{s_\Omega} + \mathbf{L_C})) \wedge ((-\mathbf{s_\Omega} + \mathbf{U_C}) \leq_d \\
&\qquad (-\mathbf{s_\Omega} + \mathbf{U_D})) \\
&\Leftrightarrow ((-\mathbf{s_\Omega} - \mathbf{L_C}) \leq_d (-\mathbf{s_\Omega} - \mathbf{L_D})) \wedge ((\mathbf{s_\Omega} - \mathbf{U_D}) \leq_d \\
&\qquad (\mathbf{s_\Omega} - \mathbf{U_C})) \\
&\Leftrightarrow (\mathbf{L}_{\overline{\eta(\mathbf{C})}} \leq_d \mathbf{L}_{\overline{\eta(\mathbf{D})}}) \wedge (\mathbf{U}_{\overline{\eta(\mathbf{D})}} \leq_d \mathbf{U}_{\overline{\eta(\mathbf{C})}}) \\
&\Leftrightarrow \overline{\eta(D)} \subseteq \overline{\eta(C)}.
\end{aligned}$$

□

D Proofs for Section 4

We start by proving Proposition 1 and Proposition 2 stated in the main text of Section 4.

Proposition 1. *Let $\mathcal{T}$ be a DL-Lite$^{\mathcal{H}}$ TBox. If $\eta \models \mathcal{T}$ then η is (weakly) TBox faithful.*

Proof. Let η be a box interpretation for a DL-Lite$^{\mathcal{H}}$ KB $\mathcal{K}$ with empty ABox and assume $\eta \models \mathcal{K}$.

Claim 9. *If $\mathcal{K}$ is a DL-Lite$^{\mathcal{H}}$ KB with an empty ABox then $\mathcal{K}$ is satisfiable.*

Proof. Let $\mathcal{I}$ be an interpretation such that for all $A \in \mathsf{N_C}$ and all $R \in \mathsf{N_R}$ we have that $A^{\mathcal{I}} = R^{\mathcal{I}} = \emptyset$. Then, trivially, $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ and $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$ where C, D are arbitrary DL-Lite$^{\mathcal{H}}$ concepts and R, S are arbitrary roles with symbols in $\mathsf{N_C}$ and $\mathsf{N_R}$. This means that, $\mathcal{I} \models \alpha$ for all (concept/role) inclusions α with symbols in $\mathsf{N_C} \cup \mathsf{N_R}$. If $\mathcal{K}$ has an empty ABox, all axioms in $\mathcal{K}$ are inclusions and thus all are satisfied by $\mathcal{I}$. In other words, $\mathcal{I} \models \mathcal{K}$, so $\mathcal{K}$ is satisfiable. □

Claim 10. *Assume $\mathcal{K}$ is a satisfiable DL-Lite$^{\mathcal{H}}$ KB with an empty ABox. If $\mathcal{K} \cup \{\alpha\}$ is unsatisfiable then α is an assertion.*

Proof. Suppose $\mathcal{K} \cup \{\alpha\}$ is unsatisfiable and $\mathcal{K}$ is a satisfiable DL-Lite$^{\mathcal{H}}$ KB with an empty ABox. By the proof of Claim 9, any interpretation $\mathcal{I}$ with $A^{\mathcal{I}} = R^{\mathcal{I}} = \emptyset$, for all $A \in \mathsf{N_C}$ and all $R \in \mathsf{N_R}$, satisfies not only $\mathcal{K}$ but also any extension of $\mathcal{K}$ with an inclusion. Thus, if $\mathcal{K} \cup \{\alpha\}$ is unsatisfiable α cannot be an inclusion. That is, α must be an assertion. □

By Claim 9, $\mathcal{K}$ is satisfiable. By Claim 10, $\mathcal{K} \cup \{\alpha\}$ is unsatisfiable only if α is an assertion. Then the lemma follows since we only require TBox faithfulness, so α in Definition 5 ranges only over inclusions, not assertions. □

Proposition 2. *Let $\mathcal{K}$ be a DL-Lite$^{\mathcal{H}}$ KB and η a box interpretation that is box consistent. If $\eta \models \mathcal{K}$ then η is (weakly) KB faithful.*

Proof. Let η be a consistent box interpretation for a DL-Lite$^{\mathcal{H}}$ KB $\mathcal{K}$ and assume $\eta \models \mathcal{K}$. Let $\mathcal{K}'$ be $\{\alpha \mid \eta \models \alpha\}$. As $\mathsf{N_I}, \mathsf{N_C}, \mathsf{N_R}$ are finite, there are finitely many axioms, so $\mathcal{K}'$ is finite. As η is box consistent, there is an interpretation $\mathcal{I}$ that satisfies $\mathcal{K}'$. Indeed, $\mathcal{I}$ can be constructed from η as follows.

- Let $\Delta^{\mathcal{I}} := \mathsf{N_I} \cup \Delta_{\mathcal{K}}$ (where $\Delta_{\mathcal{K}}$ is as in Definition 1).
- For all $a \in \mathsf{N_I}$, we have that $a^{\mathcal{I}} := a$. Moreover,
- $A^{\mathcal{I}} := \{a \in \mathsf{N_I} \mid \eta \models A(a)\} \cup \{c_D \in \Delta_{\mathcal{K}} \mid \eta \models D \sqsubseteq A\}$, for all $A \in \mathsf{N_C}$, and
- $R^{\mathcal{I}} := \{(a,b) \in \mathsf{N_I} \times \mathsf{N_I} \mid \eta \models R(a,b)\} \cup$
$\{(a, c_{\exists S}) \in \mathsf{N_I} \times \Delta_{\mathcal{K}} \mid \eta \models \exists \overline{S}(a),\ \eta \models \overline{S} \sqsubseteq R\} \cup$
$\{(c_{\exists S}, a) \in \Delta_{\mathcal{K}} \times \mathsf{N_I} \mid \eta \models \exists \overline{S}(a),\ \eta \models S \sqsubseteq R\} \cup$
$\{(c_{\exists S}, c_{\exists \overline{S}}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} \mid \eta \models S \sqsubseteq R\} \cup$
$\{(c_D, c_{\exists S}) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} | \eta \models D \sqsubseteq \exists \overline{S},\ \eta \models \overline{S} \sqsubseteq R\} \cup$
$\{(c_{\exists S}, c_D) \in \Delta_{\mathcal{K}} \times \Delta_{\mathcal{K}} \mid \eta \models D \sqsubseteq \exists \overline{S},\ \eta \models S \sqsubseteq R\}$, for all $R \in \mathsf{N_R}$.

Box Name	$\mathbf{L}[i_C]$	$\mathbf{U}[i_C]$
$I_=$	-4	-0.5
$I_\subset$	-2	-0.5
$I_\supset$	-4	2
$I_{\not\supset}$	0	2
$I_\cap$	-2	2
I_0	0	0
Ω	-4	4

Table 4: Parameters for boxes $I_=$, $I_\subset$, $I_\supset$, $I_{\not\supset}$, $I_\cap$, and Ω in dimension i_C.

Embedding Name	Value in Dimension i_C
$\mathcal{P}_\mathsf{C}$	-1
$\mathcal{P}_\mathsf{C}^{\neg}$	1
$\mathcal{B}_\mathsf{C}$	0

Table 5: Parameters for $\mathcal{P}_\mathsf{C}$ and $\mathcal{P}_\mathsf{C}^{\neg}$ in dimension i_C.

We can see that $\mathcal{I}$ is defined using η in the same way $\mathcal{I}_{\mathcal{K}}$ is defined using $\mathcal{K}$ in Definition 1 (box consistency of η ensures that $\mathcal{I}$ is well-defined). One can then employ the same argument as in Theorem 1 to show that $\eta \models \alpha$ iff $\mathcal{I} \models \alpha$. This means that $\mathcal{I} \models \mathcal{K}'$.

If $\eta \models \alpha$ then $\alpha \in \mathcal{K}'$. As $\mathcal{I} \models \mathcal{K}'$, we have that $\mathcal{I} \models \alpha$. By assumption $\eta \models \mathcal{K}$. Then, by definition, $\mathcal{K}'$ contains all axioms in $\mathcal{K}$. As $\mathcal{I} \models \mathcal{K}'$, we have that $\mathcal{I} \models \mathcal{K}$. This means that $\mathcal{I} \models \mathcal{K} \cup \{\alpha\}$. So $\mathcal{K} \cup \{\alpha\}$ is satisfiable. □

As already hinted in the main text, the proof strategy of Theorem 3 consists of first creating a mapping between finite interpretations and box interpretations and then using the canonical model for a KB to establish KB faithfulness. Definition 6 describes this mapping. The values in Tables 4, 5, 6, and 7 satisfy the constraints in Fig. 2 and Fig. 3. Any choice of values that satisfy Fig. 2 and Fig. 3 could be used.

Definition 6. *Given an interpretation $\mathcal{I}$ with finite domain, we define a box interpretation $\eta_{\mathcal{I}}$ in a d-dimensional Euclidean space, where $d = |\mathsf{N}_\mathsf{C}^{\exists}| + |\mathsf{N_R}| \cdot |\Delta^{\mathcal{I}}|$, as follows. To each concept $C \in \mathsf{N}_\mathsf{C}^{\exists}$, we associate a dimension in the vector space and denote its index by i_C. Similarly, to each pair (R, c) in $\mathsf{N_R} \times \Delta^{\mathcal{I}}$, we associate a dimension $i_{R,c}$. In our construction, we use constants defined in Tables 4, 5, 6, and 7. For each $a \in \mathsf{N_I}$, let $\eta_{\mathcal{I}}(a) := \{\mathsf{pos}_{\mathcal{I}}(a), \mathsf{bump}_{\mathcal{I}}(a)\}$ where:*

- *for every $C \in \mathsf{N}_\mathsf{C}^{\exists}$, $\mathsf{pos}_{\mathcal{I}}(a)[i_C] = \mathcal{P}_\mathsf{C}$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$, and $\mathsf{pos}_{\mathcal{I}}(a)[i_C] = \mathcal{P}_\mathsf{C}^{\neg}$ if $a^{\mathcal{I}} \notin C^{\mathcal{I}}$; also $\mathsf{bump}_{\mathcal{I}}(a)[i_C] = \mathcal{B}_\mathsf{C}$ (if $a^{\mathcal{I}} \in C^{\mathcal{I}}$ or not);*
- *for every pair (R, c) with $R \in \mathsf{N_R}$ and $c \in \Delta^{\mathcal{I}}$, $\mathsf{bump}_{\mathcal{I}}(a)[i_{R,c}] = \mathcal{B}_\mathsf{R}$ if $(c, a^{\mathcal{I}}) \in R^{\mathcal{I}}$, otherwise $\mathsf{bump}_{\mathcal{I}}(a)[i_{R,c}] = \mathcal{B}_\mathsf{R}^{\neg}$, $\mathsf{pos}_{\mathcal{I}}(a)[i_{R,c}] = \mathcal{P}_\mathsf{R}$ if $a^{\mathcal{I}} = c$, otherwise (that is, $a^{\mathcal{I}} \neq c$) $\mathsf{pos}_{\mathcal{I}}(a)[i_{R,c}] = \mathcal{P}_\mathsf{R}^{\neg}$.*

For each $D \in \mathsf{N}_\mathsf{C}^{\exists}$ with $\emptyset = D^{\mathcal{I}}$, let $\eta_{\mathcal{I}}(D)$ be the box

Box Name	$\mathbf{L}[i_{R,a}]$	$\mathbf{U}[i_{R,a}]$
I	-2	2
$\mathcal{S}_{\subset}$	-1	1
$\mathcal{S}_{\subset}^{\neg}$	-2	1
$\mathcal{B}_{\mathsf{R,I}}$	-1	1
$\mathcal{B}_{\mathsf{R,I}}^{\neg}$	-1	0
Ω	-4	4

Table 6: Parameters for boxes I, $\mathcal{S}_{\subset}$, $\mathcal{S}_{\subset}^{\neg}$, $\mathcal{B}_{\mathsf{R,I}}$, $\mathcal{B}_{\mathsf{R,I}}^{\neg}$, and Ω in dimension $i_{R,a}$.

Embedding Name	Value in Dimension $i_{R,a}$
$\mathcal{P}_{\mathsf{R}}$	$-1+\epsilon$
$\mathcal{P}_{\mathsf{R}}^{\neg}$	$1-\epsilon$
$\mathcal{B}_{\mathsf{R}}$	0
$\mathcal{B}_{\mathsf{R}}^{\neg}$	-1

Table 7: Parameters for point embeddings $\mathcal{P}_{\mathsf{R}}$, $\mathcal{P}_{\mathsf{R}}^{\neg}$, $\mathcal{B}_{\mathsf{R}}$, and $\mathcal{B}_{\mathsf{R}}^{\neg}$ in dimension $i_{R,a}$.

with all dimensions set to I_0, and, for each $D \in \mathsf{N}_{\mathsf{C}\neg}^{\exists}$ with $\emptyset \neq D^{\mathcal{I}}$, let $\eta_{\mathcal{I}}(D)$ be the box where:

- *for every $C \in \mathsf{N}_{\mathsf{C}}^{\exists}$,*
  - *$\eta_{\mathcal{I}}(D)[i_C] := I_{=}$ iff $\emptyset \neq D^{\mathcal{I}} = C^{\mathcal{I}}$,*
  - *$\eta_{\mathcal{I}}(D)[i_C] := I_{\subset}$ iff $\emptyset \neq D^{\mathcal{I}} \subset C^{\mathcal{I}}$,*
  - *$\eta_{\mathcal{I}}(D)[i_C] := I_{\supset}$ iff $D^{\mathcal{I}} \supset C^{\mathcal{I}} \neq \emptyset$,*
  - *$\eta_{\mathcal{I}}(D)[i_C] := I_{\not\supset}$ iff $\emptyset \neq C^{\mathcal{I}}$, $\emptyset \neq D^{\mathcal{I}}$, and $\emptyset = C^{\mathcal{I}} \cap D^{\mathcal{I}}$,*
  - *else $\eta_{\mathcal{I}}(D)[i_C] := I_{\cap}$;*
- *for every pair (R, c) with $R \in \mathsf{N}_{\mathsf{R}}$ and $c \in \Delta^{\mathcal{I}}$, $\eta_{\mathcal{I}}(D)[i_{R,c}] := I$.*

For each $S \in \mathsf{N}_{\mathsf{R}}$, let $\eta_{\mathcal{I}}(S)$ be the boxes $\mathsf{Head}_{\mathcal{I}}(S)$, $\mathsf{Tail}_{\mathcal{I}}(S)$, $\mathsf{Bump}_{\mathcal{I}}(S)$:

- *for every $C \in \mathsf{N}_{\mathsf{C}}^{\exists}$, $\mathsf{Head}_{\mathcal{I}}(S)[i_C] := \eta_{\mathcal{I}}(\exists S)[i_C]$, $\mathsf{Tail}_{\mathcal{I}}(S)[i_C] := \eta_{\mathcal{I}}(\exists S^{-})[i_C]$, $\mathsf{Bump}_{\mathcal{I}}(S)[i_C] := I_0$;*
- *for every pair (R, c) with $R \in \mathsf{N}_{\mathsf{R}}$ and $c \in \Delta^{\mathcal{I}}$,*
  - *$\mathsf{Head}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{S}_{\subset}$ and $\mathsf{Bump}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{B}_{\mathsf{R,I}}$ if for all $e \in \Delta^{\mathcal{I}}$ we have that $(c, e) \in S^{\mathcal{I}}$ implies $(c, e) \in R^{\mathcal{I}}$, otherwise $\mathsf{Head}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{S}_{\subset}^{\neg}$ and $\mathsf{Bump}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{B}_{\mathsf{R,I}}^{\neg}$,*
  - *$\mathsf{Tail}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{S}_{\subset}$ if for all $e \in \Delta^{\mathcal{I}}$ we have that $(e, c) \in S^{\mathcal{I}}$ implies $(c, e) \in R^{\mathcal{I}}$, otherwise $\mathsf{Tail}_{\mathcal{I}}(S)[i_{R,c}] := \mathcal{S}_{\subset}^{\neg}$.*

We first argue that $\eta_{\mathcal{I}}$ is well-defined. That is, $\eta_{\mathcal{I}}$ is indeed a box consistent interpretation. We also show that it satisfies additional properties that are used in other proofs.

Theorem 6. *Let $\mathcal{I}$ be an interpretation with finite domain, and let ϵ be an arbitrary value with $0 < \epsilon \leq \epsilon_{max}$. Then, $\eta_{\mathcal{I}}$*

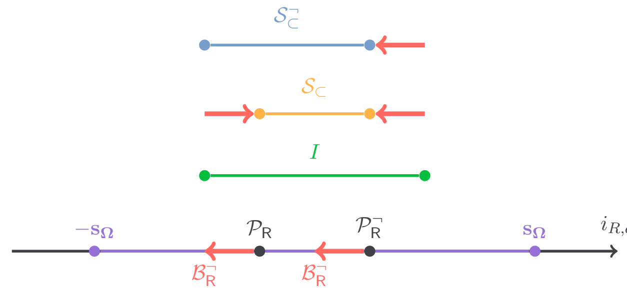

Figure 2: Visualization of the parameters of I, $\mathcal{S}_{\subset}$, $\mathcal{S}_{\subset}^{\neg}$, $\mathcal{P}_{\mathsf{R}}$, $\mathcal{P}_{\mathsf{R}}^{\neg}$, $\mathcal{B}_{\mathsf{R}}$, $\mathcal{B}_{\mathsf{R}}^{\neg}$, and Ω in dimension $i_{R,a}$.

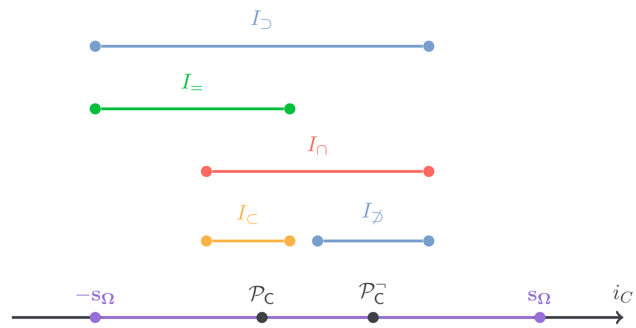

Figure 3: Visualization of the parameters of $I_{=}$, $I_{\subset}$, $I_{\supset}$, $I_{\not\supset}$, $I_{\cap}$, $\mathcal{P}_{\mathsf{C}}$, $\mathcal{P}_{\mathsf{C}}^{\neg}$ and Ω in dimension i_C.

constructed from $\mathcal{I}$, as in Definition 6, is a box interpretation that satisfies the following additional property.

- *For any concept $C \in \mathsf{N}_{\mathsf{C}}^{\exists}$, it holds that:*

$$\frac{\mathbf{L_C}[i_C] + \mathbf{U_C}[i_C]}{2} \leq -\frac{s_{\Omega}}{2},$$

which implies that $\eta_{\mathcal{I}}$ is box consistent.

Proof. Let $\mathcal{I}$ be an interpretation with finite domain and let ϵ be an arbitrary value with $0 < \epsilon \leq \epsilon_{max}$, and recall that in Section 3 we assume that ϵ_{max} is bounded by 0.5 throughout this paper. Also, $\Omega[i_C] = (-4, 4)$ (Table 6) implies $s_{\Omega} = 4$. We first prove that $\eta_{\mathcal{I}}$ as described in Definition 6 is well-defined, i.e., that it is a box interpretation (c.f. Definition 3). We start by proving the conditions related to individual names in Definition 3, namely

- each individual name $a \in \mathsf{N}_{\mathsf{I}}$ is mapped to two vectors $\eta(a) = (\mathsf{pos}(a), \mathsf{bump}(a))$, namely, a position $\mathsf{pos}(e) \in \Omega$ and a bump $\mathsf{bump}(e) \in \Omega$.

Indeed Definition 6 maps each $a \in \mathsf{N}_{\mathsf{I}}$ to two vectors: $\mathsf{pos}_{\mathcal{I}}(a)$ and $\mathsf{bump}_{\mathcal{I}}(a)$. To complete this item, we need to show that $\mathsf{pos}_{\mathcal{I}}(a), \mathsf{bump}_{\mathcal{I}}(a) \in \Omega$, which we do in Claim 11.

Claim 11. *For any individual $a \in \mathsf{N}_{\mathsf{I}}$:*

$$-\mathbf{s}_{\Omega} + \boldsymbol{\epsilon} \leq_d \mathsf{pos}_{\mathcal{I}}(a) \leq_d \mathbf{s}_{\Omega} - \boldsymbol{\epsilon}$$
$$-\mathbf{s}_{\Omega} + \boldsymbol{\epsilon} \leq_d \mathsf{bump}_{\mathcal{I}}(a) \leq_d \mathbf{s}_{\Omega} - \boldsymbol{\epsilon}$$

Proof. For any $a \in \mathsf{N_I}$ and for any dimension i with $0 \leq i \leq d$, Definition 6 assigns (i) $\mathsf{pos}_{\mathcal{I}}(a)[i]$ either to $\mathcal{P}_{\mathsf{C}}, \mathcal{P}^{-}_{\mathsf{C}}, \mathcal{P}_{\mathsf{R}}$, or $\mathcal{P}^{-}_{\mathsf{R}}$. Furthermore, by Tables 4, 5, and 7 it holds that (ii) $-s_\Omega + \epsilon_{max} \leq \mathcal{P}_{\mathsf{C}}, \mathcal{P}^{-}_{\mathsf{C}}, \mathcal{P}_{\mathsf{R}}, \mathcal{P}^{-}_{\mathsf{R}} \leq s_\Omega - \epsilon_{max}$. By (i) and (ii) for all $a \in \mathsf{N_I}$ it holds that $-\mathbf{s_\Omega} + \boldsymbol{\epsilon} \leq_d \mathsf{pos}_{\mathcal{I}}(a) \leq_d \mathbf{s_\Omega} - \boldsymbol{\epsilon}$.

For any $a \in \mathsf{N_I}$ and for any dimension i with $0 \leq i \leq d$, Definition 6 assigns (i) $\mathsf{bump}_{\mathcal{I}}(a)[i]$ either to $\mathcal{B}_{\mathsf{R}}$ or $\mathcal{B}^{-}_{\mathsf{R}}$. Also, by Tables 5, 6, and 7 it holds that (ii) $-s_\Omega + \epsilon_{max} \leq \mathcal{B}_{\mathsf{R}}, \mathcal{B}^{-}_{\mathsf{R}} \leq s_\Omega - \epsilon_{max}$. By (i) and (ii), for all $a \in \mathsf{N_I}$, $-\mathbf{s_\Omega} + \boldsymbol{\epsilon} \leq_d \mathsf{bump}_{\mathcal{I}}(a) \leq_d \mathbf{s_\Omega} - \boldsymbol{\epsilon}$. □

We now proceed with the second item of Definition 3.

- Each concept name $A \in \mathsf{N_C}$ is mapped to $\eta(A) \in \mathsf{Box}$.

To show that $\eta(A) \in \mathsf{Box}$ we need Claim 12 to hold.

Claim 12. *For any concept $C \in \mathsf{N_C}$, it holds that:*

$$\mathbf{0} \leq_d \mathbf{U_C} - \mathbf{L_C} \leq_d 2\mathbf{s_\Omega}$$

(in fact this holds for all $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$).

Proof. For any concept $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$ and for any dimension i with $0 \leq i \leq d$, Definition 6 assigns $\eta_{\mathcal{I}}(C)[i]$ either to $I_0, I_=, I_\subset, I_\supset, I_{\not\supset}, I_\cap$, or I. Then, by Tables 4 and 6, for $(\mathbf{L_C}[i], \mathbf{U_C}[i]) := \eta_{\mathcal{I}}(C)[i]$ it holds that $0 \leq \mathbf{U_C}[i] - \mathbf{L_C}[i] \leq 2s_\Omega$. Thus, $\mathbf{0} \leq_d \mathbf{U_C} - \mathbf{L_C} \leq_d 2\mathbf{s_\Omega}$. □

The third item of Definition 3 is:

- each role name $R \in \mathsf{N_R}$ is mapped to three boxes $\eta(R) = (\mathsf{Head}(R), \mathsf{Tail}(R), \mathsf{Bump}(R))$, which we call R's head $\mathsf{Head}(R)$, tail $\mathsf{Tail}(R)$ and bump box $\mathsf{Bump}(R)$.

We show this in Claim 13.

Claim 13. *For any role $R \in \mathsf{N_R}$, it holds that:*

$$\mathbf{0} \leq_d \mathbf{U}^X_{\mathbf{R}} - \mathbf{L}^X_{\mathbf{R}} \leq_d 2\mathbf{s_\Omega} \text{ with } X \in \{\mathbf{H}, \mathbf{T}, \mathbf{B}\}.$$

Proof. For any $R \in \mathsf{N_R}$ and for any dimension i with $0 \leq i \leq d$, Definition 6 assigns $\mathsf{Head}_{\mathcal{I}}(R)[i]$ and $\mathsf{Tail}_{\mathcal{I}}(R)[i]$ either to $\mathcal{S}_\subset, \mathcal{S}^{-}_\subset, I_0, I_=, I_\subset, I_\supset, I_{\not\supset}$, or $I_\cap$. Then, by Tables 4 and 6, for $(\mathbf{L}^{\mathbf{H}}_{\mathbf{R}}[i], \mathbf{U}^{\mathbf{H}}_{\mathbf{R}}[i]) := \mathsf{Head}_{\mathcal{I}}(R)[i]$ it holds that $0 \leq \mathbf{U}^{\mathbf{H}}_{\mathbf{R}}[i] - \mathbf{L}^{\mathbf{H}}_{\mathbf{R}}[i] \leq 2s_\Omega$, and for $(\mathbf{L}^{\mathbf{T}}_{\mathbf{R}}[i], \mathbf{U}^{\mathbf{T}}_{\mathbf{R}}[i]) := \mathsf{Tail}_{\mathcal{I}}(R)[i]$ it holds that $0 \leq \mathbf{U}^{\mathbf{T}}_{\mathbf{R}}[i] - \mathbf{L}^{\mathbf{T}}_{\mathbf{R}}[i] \leq 2s_\Omega$.

For any $R \in \mathsf{N_R}$ and for any dimension i with $0 \leq i \leq d$, Definition 6 assigns $\mathsf{bump}_{\mathcal{I}}(R)[i]$ either to $I_0, \mathcal{B}_{\mathsf{R,I}}$, or $\mathcal{B}^{-}_{\mathsf{R,I}}$. Then, by Tables 4 and 6, for $(\mathbf{L}^{\mathbf{B}}_{\mathbf{R}}[i], \mathbf{U}^{\mathbf{B}}_{\mathbf{R}}[i]) := \mathsf{bump}_{\mathcal{I}}(R)[i]$ it holds that $0 \leq \mathbf{U}^{\mathbf{B}}_{\mathbf{R}}[i] - \mathbf{L}^{\mathbf{B}}_{\mathbf{R}}[i] \leq 2s_\Omega$. □

The extension of $\eta_{\mathcal{I}}$ to arbitrary DL-Lite$^{\mathcal{H}}$ concept and role expressions is as in Definition 3, though, for presentation purposes we did not explicitly define the head and tail boxes for the dimension i_C (this was defined in terms of $\eta_{\mathcal{I}}(\exists S)$ and $\eta_{\mathcal{I}}(\exists S^-)$). So here we need to argue that $\eta_{\mathcal{I}}(\exists S)$ is equal to

$$\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L}^{\mathbf{H}}_{\mathbf{S}} - \mathbf{U}^{\mathbf{B}}_{\mathbf{S}} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U}^{\mathbf{H}}_{\mathbf{S}} - \mathbf{L}^{\mathbf{B}}_{\mathbf{S}} - \boldsymbol{\epsilon}\}$$

(recall that $\mathbf{L}_{\exists\mathbf{S}} = \mathbf{L}^{\mathbf{H}}_{\mathbf{S}} - \mathbf{U}^{\mathbf{B}}_{\mathbf{S}}$ and $\mathbf{U}_{\exists\mathbf{S}} = \mathbf{U}^{\mathbf{H}}_{\mathbf{S}} - \mathbf{L}^{\mathbf{B}}_{\mathbf{S}}$ see Definition 3). Indeed, by the values in Table 6, for the dimensions of the form $i_{R,c}$ we have that $\mathbf{L}^{\mathbf{H}}_{\mathbf{S}}[i_{R,c}] - \mathbf{U}^{\mathbf{B}}_{\mathbf{S}}[i_{R,c}] = -2$ and $\mathbf{U}^{\mathbf{H}}_{\mathbf{S}}[i_{R,c}] - \mathbf{L}^{\mathbf{B}}_{\mathbf{S}}[i_{R,c}] = 2$, which correspond to I, the value of $\eta_{\mathcal{I}}(\exists S)[i_{R,c}]$. For the dimensions of the form i_C, $\mathbf{L}^{\mathbf{H}}_{\mathbf{S}}[i_C] - \mathbf{U}^{\mathbf{B}}_{\mathbf{S}}[i_C] = \mathbf{L}^{\mathbf{H}}_{\mathbf{S}}[i_C]$ and $\mathbf{U}^{\mathbf{H}}_{\mathbf{S}}[i_C] - \mathbf{L}^{\mathbf{B}}_{\mathbf{S}}[i_C] = \mathbf{U}^{\mathbf{H}}_{\mathbf{S}}[i_C]$ since in this case $\mathsf{Bump}_{\mathcal{I}}(S)[i_C] := I_0$, which is $(0, 0)$, by the values in Table 4. This is as required as in this case we have $\mathsf{Head}_{\mathcal{I}}(S)[i_C] = \eta_{\mathcal{I}}(\exists S)[i_C]$. The argument that $\eta_{\mathcal{I}}(\exists S^-)$ is equal to

$$\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L}^{\mathbf{T}}_{\mathbf{S}} - \mathbf{U}^{\mathbf{B}}_{\mathbf{S}} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U}^{\mathbf{T}}_{\mathbf{S}} - \mathbf{L}^{\mathbf{B}}_{\mathbf{S}} - \boldsymbol{\epsilon}\}$$

is similar. It remains to argue the following.

Claim 14. *For any concept $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$, it holds that:*

$$\frac{\mathbf{L_C}[i_C] + \mathbf{U_C}[i_C]}{2} \leq -\frac{s_\Omega}{2} \tag{9}$$

and that this implies $\eta_{\mathcal{I}}$ is box consistent.

Proof. For any concept $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$ and dimension i_C, (i) Definition 6 assigns $\eta_{\mathcal{I}}(C)[i_C] = I_=$. Then, by Table 4 it holds that (ii) the lower bound of $I_=$ is set to -4 and the upper bound of $I_=$ is set to -0.5. By (i) and (ii), we have that (iii) $\frac{\mathbf{L_C}[i_C] + \mathbf{U_C}[i_C]}{2} = -2.25 \leq -s_\Omega/2 = -2$. From (i), (ii), and (iii), for any concept $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$, it holds that $\frac{\mathbf{L_C}[i_C] + \mathbf{U_C}[i_C]}{2} \leq -\frac{s_\Omega}{2}$, proving (Equation (9)).

Finally, we briefly check that (9) guarantees box consistency. Let C be an arbitrary concept $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$ and assume that (9) holds. Next, recall that $\mathbf{L}_{\neg\mathbf{C}} = -\mathbf{s_\Omega} - \mathbf{L_C}$ holds (see Section 5.1). Now we have that:

$$\mathbf{U_C}[i_C] \leq -s_\Omega - \mathbf{L_C}[i_C] = \mathbf{L}_{\neg\mathbf{C}}[i_C] \tag{10}$$

In particular,

$$\mathbf{U_C}[i_C] - \epsilon < \mathbf{L}_{\neg\mathbf{C}}[i_C] + \epsilon,$$

which means that the boxes defined by C and $\neg C$ do not intersect in dimension i_C. This implies that $\eta_{\mathcal{I}}(C) \cap \eta_{\mathcal{I}}(\neg C) \neq \emptyset$, proving that $\eta_{\mathcal{I}}$ is box consistent. □

This finishes the proof of this theorem. □

Fig. 3 and Fig. 2 help to visualize how the constants defined in Tables 4, 5, 6, and 7 relate to each other. In the next lemmas, we argue that $\mathcal{I} \models \alpha$ iff $\eta_{\mathcal{I}} \models \alpha$ when α is a concept assertion, role assertion, concept inclusion, or role inclusion.

Lemma 1. *Given an interpretation $\mathcal{I}$ with finite domain, let $\eta_{\mathcal{I}}$ be as in Definition 6. For all $a \in \mathsf{N_I}$ and all concepts $D \in \mathsf{N}^{\exists}_{\mathsf{C}}$, $\mathcal{I} \models D(a)$ iff $\eta_{\mathcal{I}} \models D(a)$.*

Proof. We need to show that $a^{\mathcal{I}} \in D^{\mathcal{I}}$ iff $\mathsf{pos}_{\mathcal{I}}(a) \in \eta_{\mathcal{I}}(D)$. ($\Rightarrow$) Suppose $a^{\mathcal{I}} \in D^{\mathcal{I}}$. To show that $\mathsf{pos}_{\mathcal{I}}(a) \in \eta_{\mathcal{I}}(D)$, we argue that $\mathsf{pos}_{\mathcal{I}}(a)[j] \in \eta_{\mathcal{I}}(D)[j]$, for every dimension $1 \leq j \leq d$. We make a case distinction with $C \in \mathsf{N}^{\exists}_{\mathsf{C}}$.

- **Dimension i_C with $a^\mathcal{I} \in C^\mathcal{I}$.** In this case, (i) $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}_\mathsf{C}$. By the parameters of the boxes of dimension i_C (see Tables 4 and 5, and Figure 3), it holds that (ii) $\mathcal{P}_\mathsf{C} \in I_= \cap I_\supset \cap I_\subset \cap I_\cap$. By Definition 6, $\eta_\mathcal{I}(D)[i_C]$ is either equal to $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$ (since by the ($\Rightarrow$) assumption $D^\mathcal{I} \neq \emptyset$). By the ($\Rightarrow$) assumption, $a^\mathcal{I} \in D^\mathcal{I}$ and by the assumption of this case $a^\mathcal{I} \in C^\mathcal{I}$. By Definition 6, we have that $\eta_\mathcal{I}(D)[i_C]$ is not $I_{\not\supset}$ because $a^\mathcal{I} \in C^\mathcal{I} \cap D^\mathcal{I}$; so it holds that (iii) $\eta_\mathcal{I}(D)[i_C] \neq I_{\not\supset}$. By (i)-(iii) it holds that $\mathsf{pos}_\mathcal{I}(a)[i_C] \in \eta_\mathcal{I}(D)[i_C]$.
- **Dimension i_C with $a^\mathcal{I} \notin C^\mathcal{I}$.** In this case, (i) $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}^-_\mathsf{C}$. By the parameters of the boxes of dimension i_C (see Tables 4 and 5, and Figure 3), it holds that (ii) $\mathcal{P}^-_\mathsf{C} \in I_{\not\supset} \cap I_\supset \cap I_\cap$. By Definition 6, $\eta_\mathcal{I}(D)[i_C]$ is either equal to $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$, (since by the ($\Rightarrow$) assumption $D^\mathcal{I} \neq \emptyset$). By the ($\Rightarrow$) assumption, $a^\mathcal{I} \in D^\mathcal{I}$ and by the assumption of this case $a^\mathcal{I} \notin C^\mathcal{I}$. Also, by Definition 6 we have that $\eta_\mathcal{I}(D)[i_C]$ is neither $I_=$ nor $I_\subset$ because the former requires $C^\mathcal{I} = D^\mathcal{I}$, the latter requires $D^\mathcal{I} \subset C^\mathcal{I}$, but in the case we consider here we have $a^\mathcal{I} \in D^\mathcal{I} \setminus C^\mathcal{I}$; so (iii) $\eta_\mathcal{I}(D)[i_C] \neq I_=$ and (iv) $\eta_\mathcal{I}(D)[i_C] \neq I_\subset$. By (i)-(iv) it holds that $\mathsf{pos}_\mathcal{I}(a)[i_C] \in \eta_\mathcal{I}(D)[i_C]$.

Finally, consider the dimensions of the form $i_{R,c}$, where $R \in \mathsf{N_R}$ and $c \in \Delta^\mathcal{I}$. In this case (i) $\eta_\mathcal{I}(D)[i_{R,c}] = I$ for any $R \in \mathsf{N_R}$ and $c \in \Delta^\mathcal{I}$. Furthermore, for any individual name in $\mathsf{N_I}$, in particular a, we have that (ii) $\mathsf{pos}_\mathcal{I}(a)[i_{R,c}] = \mathcal{P}_\mathsf{R}$ or $\mathsf{pos}_\mathcal{I}(a)[i_{R,c}] = \mathcal{P}^-_\mathsf{R}$. By the parameters of I, $\mathcal{P}_\mathsf{R}$, and $\mathcal{P}^-_\mathsf{R}$ (see Tables 6 and 7, and Figure 2), it holds that (iii) $\mathcal{P}_\mathsf{R} \in I$ and $\mathcal{P}^-_\mathsf{R} \in I$. By (i)-(iii) it holds that $\mathsf{pos}_\mathcal{I}(a)[i_{R,c}] \in \eta_\mathcal{I}(D)[i_{R,c}]$ provided that $\epsilon \leq \epsilon_{max}$ (see Table 7 and Section 3). We have shown that $\mathsf{pos}_\mathcal{I}(a)[j] \in \eta_\mathcal{I}(D)[j]$ for any dimension $1 \leq j \leq d$. Thus, we have shown that $\mathsf{pos}_\mathcal{I}(a) \in \eta_\mathcal{I}(D)$ if $a^\mathcal{I} \in D^\mathcal{I}$.

($\Leftarrow$) Now suppose $\mathsf{pos}_\mathcal{I}(a) \in \eta_\mathcal{I}(D)$. This means, for dimension i_D in particular, that (i) $\mathsf{pos}_\mathcal{I}(a)[i_D] \in \eta_\mathcal{I}(D)[i_D]$. By the construction of dimension i_D of $\eta_\mathcal{I}$ it holds that (ii) $\eta_\mathcal{I}(D)[i_D] = I_=$. From (i) and (ii), $\mathsf{pos}_\mathcal{I}(a)[i_D] = \mathcal{P}_\mathsf{C}$. By the construction of dimension i_D $\mathsf{pos}_\mathcal{I}(a)[i_D] = \mathcal{P}_\mathsf{C}$ can only be if $a^\mathcal{I} \in D^\mathcal{I}$. □

Lemma 2. *Given an interpretation $\mathcal{I}$ with finite domain, let $\eta_\mathcal{I}$ be as in Definition 6. For all $a, b \in \mathsf{N_I}$ and all $R \in \mathsf{N_R}$, $\mathcal{I} \models R(a,b)$ iff $\eta_\mathcal{I} \models R(a,b)$.*

Proof. ($\Rightarrow$) Assume $\mathcal{I} \models R(a,b)$. We start by showing that $\mathsf{pos}_\mathcal{I}(a) + \mathsf{bump}_\mathcal{I}(b) \in \mathsf{Head}_\mathcal{I}(R)$ by showing that this holds in every dimension. For this we make a case distinction, where $C \in \mathsf{N}^\exists_\mathsf{C}$, $S \in \mathsf{N_R}$, and $c \in \Delta^\mathcal{I}$.

- **Dimension i_C.** We want to show that $\mathsf{pos}_\mathcal{I}(a)[i_C] + \mathsf{bump}_\mathcal{I}(b)[i_C] \in \mathsf{Head}_\mathcal{I}(R)[i_C]$. By Definition 6, $\mathsf{Head}_\mathcal{I}(R)[i_C] = \eta_\mathcal{I}(\exists R)[i_C]$ can be either equal to $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$ (since by the ($\Rightarrow$) assumption $\mathcal{I} \models R(a,b)$, so $(\exists R)^\mathcal{I} \neq \emptyset$). By Definition 6, $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}_\mathsf{C}$ if $a^\mathcal{I} \in C^\mathcal{I}$, and $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}^-_\mathsf{C}$ if $a^\mathcal{I} \notin C^\mathcal{I}$; also $\mathsf{bump}_\mathcal{I}(a)[i_C] = \mathcal{B}_\mathsf{C}$ (if $a^\mathcal{I} \in C^\mathcal{I}$ or not). By the values in Tables 4 and 5, we have that all possible values for $\mathsf{pos}_\mathcal{I}(a)[i_C] + \mathsf{bump}_\mathcal{I}(b)[i_C]$ are in $I_\supset \cap I_\cap$. So we need to consider three cases, namely, (i) when $\mathsf{Head}_\mathcal{I}(R)[i_C] = I_=$, (ii) when $\mathsf{Head}_\mathcal{I}(R)[i_C] = I_\subset$, and (iii) when $\mathsf{Head}_\mathcal{I}(R)[i_C] = I_{\not\supset}$. By the ($\Rightarrow$) assumption, $\mathcal{I} \models R(a,b)$, so $a^\mathcal{I} \in (\exists R)^\mathcal{I}$. In case (i) and in case (ii), by Definition 6, we have that $(\exists R)^\mathcal{I} = C^\mathcal{I}$ and $(\exists R)^\mathcal{I} \subset C^\mathcal{I}$, respectively. So, in both cases $a^\mathcal{I} \in C^\mathcal{I}$ and, by Definition 6, $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}_\mathsf{C}$. By the values in Tables 4 and 5, $\mathsf{pos}_\mathcal{I}(a)[i_C] + \mathsf{bump}_\mathcal{I}(b)[i_C] \in I_= \cap I_\subset$. In case (iii), $(\exists R)^\mathcal{I} \cap C^\mathcal{I} = \emptyset$. So $a^\mathcal{I} \notin C^\mathcal{I}$ and, by Definition 6, $\mathsf{pos}_\mathcal{I}(a)[i_C] = \mathcal{P}^-_\mathsf{C}$. By the values in Tables 4 and 5, $\mathsf{pos}_\mathcal{I}(a)[i_C] + \mathsf{bump}_\mathcal{I}(b)[i_C] \in I_{\not\supset}$.
- **Dimension $i_{S,c}$ with $a^\mathcal{I} = c$ and for all $e \in \Delta^\mathcal{I}$, $(c,e) \in R^\mathcal{I}$ implies $(c,e) \in S^\mathcal{I}$.** By Definition 6, when $a^\mathcal{I} = c$, we have that $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] = \mathcal{P}_\mathsf{R}$ and $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] = \mathcal{B}_\mathsf{R}$. Also, in this case $\mathsf{Head}_\mathcal{I}(R)[i_{S,c}]$ is $\mathcal{S}_\subset$. By the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(b)[i_{S,c}] \in \mathsf{Head}_\mathcal{I}(R)[i_{S,c}]$.
- **Dimension $i_{S,c}$ with $a^\mathcal{I} = c$ and there is $e \in \Delta^\mathcal{I}$, with $(c,e) \in R^\mathcal{I}$ but $(c,e) \notin S^\mathcal{I}$.** By Definition 6, when $a^\mathcal{I} = c$, we have that $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] = \mathcal{P}_\mathsf{R}$ and $\mathsf{Head}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}^-_\subset$. Also, $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ or $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] = \mathcal{B}^-_\mathsf{R}$. In both cases, by the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(b)[i_{S,c}] \in \mathsf{Head}_\mathcal{I}(R)[i_{S,c}]$.
- **Dimension $i_{S,c}$ with $a^\mathcal{I} \neq c$.** In this case, by Definition 6, $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] = \mathcal{P}^-_\mathsf{R}$. It can be that $\mathsf{Head}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}_\subset$ or $\mathsf{Head}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}^-_\subset$. Also, it can be that $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] = \mathcal{B}^-_\mathsf{R}$ or $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ (depending on $\mathcal{I}$) but in all cases, by the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(a)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(b)[i_{S,c}] \in \mathsf{Head}_\mathcal{I}(R)[i_{S,c}]$, as required.

We now show that $\mathsf{pos}_\mathcal{I}(b) + \mathsf{bump}_\mathcal{I}(a) \in \mathsf{Tail}_\mathcal{I}(R)$ by showing that this holds in every dimension. We make a similar case distinction as in the argument above.

- **Dimension i_C.** We want to show that $\mathsf{pos}_\mathcal{I}(b)[i_C] + \mathsf{bump}_\mathcal{I}(a)[i_C] \in \mathsf{Tail}_\mathcal{I}(R)[i_C]$. By Definition 6, $\mathsf{Tail}_\mathcal{I}(R)[i_C] = \eta_\mathcal{I}(\exists R^-)[i_C]$ can be either equal to $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$ (since by the ($\Rightarrow$) assumption $\mathcal{I} \models R(a,b)$, so $(\exists R^-)^\mathcal{I} \neq \emptyset$). By Definition 6, $\mathsf{pos}_\mathcal{I}(b)[i_C] = \mathcal{P}_\mathsf{C}$ if $b^\mathcal{I} \in C^\mathcal{I}$, and $\mathsf{pos}_\mathcal{I}(b)[i_C] = \mathcal{P}^-_\mathsf{C}$ if $b^\mathcal{I} \notin C^\mathcal{I}$; also $\mathsf{bump}_\mathcal{I}(b)[i_C] = \mathcal{B}_\mathsf{C}$ (if $b^\mathcal{I} \in C^\mathcal{I}$ or not). By the values in Tables 4 and 5, we have that all possible values for $\mathsf{pos}_\mathcal{I}(b)[i_C] + \mathsf{bump}_\mathcal{I}(a)[i_C]$ are in $I_\supset \cap I_\cap$. So we need to consider three cases, namely, (i) when $\mathsf{Tail}_\mathcal{I}(R)[i_C] = I_=$, (ii) when $\mathsf{Tail}_\mathcal{I}(R)[i_C] = I_\subset$, and (iii) when $\mathsf{Tail}_\mathcal{I}(R)[i_C] = I_{\not\supset}$. By the ($\Rightarrow$) assumption, $\mathcal{I} \models R(a,b)$, so $b^\mathcal{I} \in (\exists R^-)^\mathcal{I}$. In case (i) and in case (ii), by Definition 6, we have that $(\exists R^-)^\mathcal{I} = C^\mathcal{I}$ and $(\exists R^-)^\mathcal{I} \subset C^\mathcal{I}$, respectively. So, in both cases $b^\mathcal{I} \in C^\mathcal{I}$ and, by Definition 6, $\mathsf{pos}_\mathcal{I}(b)[i_C] = \mathcal{P}_\mathsf{C}$. By the values in Tables 4 and 5, $\mathsf{pos}_\mathcal{I}(b)[i_C] + \mathsf{bump}_\mathcal{I}(a)[i_C] \in I_= \cap I_\subset$. In case (iii), $(\exists R^-)^\mathcal{I} \cap C^\mathcal{I} = \emptyset$. So $b^\mathcal{I} \notin C^\mathcal{I}$ and, by Definition 6, $\mathsf{pos}_\mathcal{I}(b)[i_C] = \mathcal{P}^-_\mathsf{C}$. By the values in Tables 4 and 5, $\mathsf{pos}_\mathcal{I}(b)[i_C] + \mathsf{bump}_\mathcal{I}(a)[i_C] \in I_{\not\supset}$.

- **Dimension** $i_{S,c}$ **with** $b^\mathcal{I} = c$ **and for all** $e \in \Delta^\mathcal{I}$**,** $(e,c) \in R^\mathcal{I}$ **implies** $(c,e) \in S^\mathcal{I}$**.** By Definition 6, when $b^\mathcal{I} = c$, we have that $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] = \mathcal{P}_\mathsf{R}$ and $\mathsf{Tail}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}_\subset$. We also have that $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ because $\mathcal{I} \models R(a,b)$ and in our particular case this also implies $\mathcal{I} \models R(b,a)$. By the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(a)[i_{S,c}] \in \mathsf{Tail}_\mathcal{I}(R)[i_{S,c}]$.

- **Dimension** $i_{S,c}$ **with** $b^\mathcal{I} = c$ **and there is** $e \in \Delta^\mathcal{I}$**, with**$(e,c) \in R^\mathcal{I}$ **but** $(c,e) \notin S^\mathcal{I}$**.** By Definition 6, when $b^\mathcal{I} = c$, we have that $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] = \mathcal{P}_\mathsf{R}$ and $\mathsf{Tail}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}_\subset^\neg$. Also, $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ or $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}^\neg$. In both cases, by the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(a)[i_{S,c}] \in \mathsf{Tail}_\mathcal{I}(R)[i_{S,c}]$.

- **Dimension** $i_{S,c}$ **with** $b^\mathcal{I} \neq c$**.** In this case, by Definition 6, $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] = \mathcal{P}_\mathsf{R}^\neg$. It can be that $\mathsf{Tail}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}_\subset$ or $\mathsf{Tail}_\mathcal{I}(R)[i_{S,c}] = \mathcal{S}_\subset^\neg$. Also, it can be that $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}^\neg$ or $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ (depending on $\mathcal{I}$) but in all cases, by the values in Tables 6 and 7, $\mathsf{pos}_\mathcal{I}(b)[i_{S,c}] + \mathsf{bump}_\mathcal{I}(a)[i_{S,c}] \in \mathsf{Tail}_\mathcal{I}(R)[i_{S,c}]$.

It remains to show that $\mathsf{bump}_\mathcal{I}(a) \in \mathsf{Bump}_\mathcal{I}(R)$ and $\mathsf{bump}_\mathcal{I}(b) \in \mathsf{Bump}_\mathcal{I}(R)$. We again make a case distinction.

- **Dimension** i_C**.** By Definition 6, we have that $\mathsf{Bump}_\mathcal{I}(R)[i_C] = I_0$ and $\mathsf{bump}_\mathcal{I}(a)[i_C] = \mathsf{bump}_\mathcal{I}(b)[i_C] = \mathcal{B}_\mathsf{C}$. By the values in Tables 4 and 5, $\mathsf{bump}_\mathcal{I}(a)[i_C] = \mathsf{bump}_\mathcal{I}(b)[i_C] \in \mathsf{Bump}_\mathcal{I}(R)[i_C]$.

- **Dimension** $i_{S,c}$**.** By Definition 6, we have that $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}$ or $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] = \mathcal{B}_\mathsf{R}^\neg$ and the same for $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}]$. Also, $\mathsf{Bump}_\mathcal{I}(R)[i_{S,c}] = \mathcal{B}_{\mathsf{R},\mathsf{I}}$ or $\mathsf{Bump}_\mathcal{I}(R)[i_{S,c}] = \mathcal{B}_{\mathsf{R},\mathsf{I}}^\neg$. In all cases, by the values in Tables 6 and 7, $\mathsf{bump}_\mathcal{I}(a)[i_{S,c}] \in \mathsf{Bump}_\mathcal{I}(R)[i_{S,c}]$ and $\mathsf{bump}_\mathcal{I}(b)[i_{S,c}] \in \mathsf{Bump}_\mathcal{I}(R)[i_{S,c}]$.

We have thus shown that

$$\begin{aligned}\mathsf{pos}_\mathcal{I}(a) + \mathsf{bump}_\mathcal{I}(b) &\in \mathsf{Head}_\mathcal{I}(R)\\ \mathsf{pos}_\mathcal{I}(b) + \mathsf{bump}_\mathcal{I}(a) &\in \mathsf{Tail}_\mathcal{I}(R)\\ \mathsf{bump}_\mathcal{I}(a) &\in \mathsf{Bump}_\mathcal{I}(R)\\ \mathsf{bump}_\mathcal{I}(b) &\in \mathsf{Bump}_\mathcal{I}(R).\end{aligned}$$

Then, by Definition 4, we have that $\eta_\mathcal{I} \models R(a,b)$.

($\Leftarrow$) Assume $\eta_\mathcal{I} \models R(a,b)$. Let $c \in \Delta^\mathcal{I}$ be the element such that $a^\mathcal{I} = c$. By Definition 6, $\mathsf{pos}_\mathcal{I}(a)[i_{R,c}] = \mathcal{P}_\mathsf{R}$ and $\mathsf{Head}_\mathcal{I}(R)[i_{R,c}] := \mathcal{S}_\subset$. Also, either $\mathsf{bump}_\mathcal{I}(b)[i_{R,c}] = \mathcal{B}_\mathsf{R}$ or $\mathsf{bump}_\mathcal{I}(b)[i_{R,c}] = \mathcal{B}_\mathsf{R}^\neg$. By Definition 4, $\mathsf{pos}_\mathcal{I}(a) + \mathsf{bump}_\mathcal{I}(b) \in \mathsf{Head}_\mathcal{I}(R)$. This means, in particular, that for the dimension $i_{R,c}$

$$\mathsf{pos}_\mathcal{I}(a)[i_{R,c}] + \mathsf{bump}_\mathcal{I}(b)[i_{R,c}] \in \mathsf{Head}_\mathcal{I}(R)[i_{R,c}].$$

The only possible value of $\mathsf{bump}_\mathcal{I}(b)[i_{R,c}]$ that satisfies the conditions above is $\mathcal{B}_\mathsf{R}$, which is assigned to $\mathsf{bump}_\mathcal{I}(b)[i_{R,c}]$ iff $(c, b^\mathcal{I}) \in R^\mathcal{I}$. Since $a^\mathcal{I} = c$, we have that $(a^\mathcal{I}, b^\mathcal{I}) \in R^\mathcal{I}$. In other words, $\mathcal{I} \models R(a,b)$. □

Lemma 3. *Given an interpretation* $\mathcal{I}$ *with finite domain, let* $\eta_\mathcal{I}$ *be as in Definition 6. For all DL-Lite*$^\mathcal{H}$ *CIs* $C \sqsubseteq D$, $\mathcal{I} \models C \sqsubseteq D$ *iff* $\eta_\mathcal{I} \models C \sqsubseteq D$.

Proof. We need to show that $C^\mathcal{I} \subseteq D^\mathcal{I}$ iff $\eta_\mathcal{I}(C) \subseteq \eta_\mathcal{I}(D)$. ($\Rightarrow$) Suppose $C^\mathcal{I} \subseteq D^\mathcal{I}$. To show that $\eta_\mathcal{I}(C) \subseteq \eta_\mathcal{I}(D)$, we argue that $\eta_\mathcal{I}(C)[j] \subseteq \eta_\mathcal{I}(D)[j]$, for every dimension $1 \leq j \leq d$. First assume $\emptyset = D^\mathcal{I}$. Then, by the ($\Rightarrow$) assumption, $\emptyset = C^\mathcal{I}$. By the box definition (see Section 3), we have that $\eta_\mathcal{I}(C) = \eta_\mathcal{I}(D) = \emptyset$ (when $\epsilon > 0$, which is assumed to be the case in this work). So $\eta_\mathcal{I}(C) \subseteq \eta_\mathcal{I}(D)$. We now make a case distinction with $\emptyset \neq D^\mathcal{I}$ and $E \in \mathsf{N}_\mathsf{C}^\exists$.

- **Dimension** i_E **with** $\emptyset \neq D^\mathcal{I} = E^\mathcal{I}$**.** By Definition 6, $\eta_\mathcal{I}(E)[i_E] = I_=$. By the assumption in this case (i) $D^\mathcal{I} = E^\mathcal{I}$, so (ii) $\eta_\mathcal{I}(D)[i_E] = I_=$ by Definition 6. By the ($\Rightarrow$) assumption, $C^\mathcal{I} \subseteq D^\mathcal{I}$. We have two cases.
  - If $C^\mathcal{I} = D^\mathcal{I}$ then, by (i), $C^\mathcal{I} = E^\mathcal{I}$. So, by Definition 6, we have that $\eta_\mathcal{I}(C)[i_E] = I_=$, which implies $\eta_\mathcal{I}(C)[i_E] = \eta_\mathcal{I}(D)[i_E]$ and then $\eta_\mathcal{I}(C)[i_E] \subseteq \eta_\mathcal{I}(D)[i_E]$, as required.
  - Otherwise, $C^\mathcal{I} \subset D^\mathcal{I}$. Then, by (i), $C^\mathcal{I} \subset E^\mathcal{I}$. If $C^\mathcal{I} \neq \emptyset$ then by Definition 6, $\eta_\mathcal{I}(C)[i_E] = I_\subset$. By (ii), $\eta_\mathcal{I}(D)[i_E] = I_=$. Since $I_\subset \subseteq I_=$ (see Table 4 and Figure 3), $\eta_\mathcal{I}(C)[i_E] \subseteq \eta_\mathcal{I}(D)[i_E]$. Otherwise (when $C^\mathcal{I} = \emptyset$) then $\eta_\mathcal{I}(C) = \emptyset$ and we are done.

- **Dimension** i_E **with** $\emptyset \neq D^\mathcal{I} \subset E^\mathcal{I}$**.** By Definition 6 and the assumption in this case (i) $\eta_\mathcal{I}(D)[i_E] = I_\subset$. By the ($\Rightarrow$) assumption $C^\mathcal{I} \subseteq D^\mathcal{I}$ and by the assumption in this case $D^\mathcal{I} \subset E^\mathcal{I}$, so $C^\mathcal{I} \subset E^\mathcal{I}$. If $C^\mathcal{I} \neq \emptyset$ then by Definition 6, (ii) $\eta_\mathcal{I}(C)[i_E] = I_\subset$. Then, by (i) and (ii) it holds that $\eta_\mathcal{I}(C)[i_E] = \eta_\mathcal{I}(D)[i_E]$ so $\eta_\mathcal{I}(C)[i_E] \subseteq \eta_\mathcal{I}(D)[i_E]$. Otherwise (when $C^\mathcal{I} = \emptyset$), $\eta_\mathcal{I}(C) = \emptyset$ and we are done.

- **Dimension** i_E **with** $D^\mathcal{I} \supset E^\mathcal{I} \neq \emptyset$**.** By Definition 6 and the assumption in this case (i) $\eta_\mathcal{I}(D)[i_E] = I_\supset$. By Definition 6 $\eta_\mathcal{I}(C)[i_E]$ is either equal to (ii) $\emptyset$, $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$. By the chosen parameters of $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, and $I_{\not\supset}$ (see Table 4 and Figure 3) it holds that (iii) $I_= \cup I_\supset \cup I_\subset \cup I_\cap \cup I_{\not\supset} \subseteq I_\supset$. Then, by (i)-(iii), any possible value of $\eta_\mathcal{I}(C)[i_E]$ leads to $\eta_\mathcal{I}(C)[i_E] \subseteq \eta_\mathcal{I}(D)[i_E]$, as required.

- **Dimension** i_E **with** $\emptyset \neq E^\mathcal{I}$**,** $\emptyset \neq D^\mathcal{I}$**, and** $\emptyset = E^\mathcal{I} \cap D^\mathcal{I}$**.** By Definition 6 and the assumption in this case (i) $\eta_\mathcal{I}(D)[i_E] = I_{\not\supset}$. By the ($\Rightarrow$) assumption $C^\mathcal{I} \subseteq D^\mathcal{I}$, so $\emptyset = E^\mathcal{I} \cap C^\mathcal{I}$. If $\emptyset \neq C^\mathcal{I}$ then by Definition 6 (ii) $\eta_\mathcal{I}(C)[i_E] = I_{\not\supset}$. By (i) and (ii) $\eta_\mathcal{I}(C)[i_E] = \eta_\mathcal{I}(D)[i_E]$, so $\eta_\mathcal{I}(C)[i_E] \subseteq \eta_\mathcal{I}(D)[i_E]$. The case where $C^\mathcal{I} = \emptyset$ is as argued above.

- **Dimension** i_E **with none of the above.** By Definition 6 and the assumption in this case (i) $\eta_\mathcal{I}(D)[i_E] = I_\cap$. Again by Definition 6, $\eta_\mathcal{I}(C)[i_E]$ is either equal to (ii) $\emptyset$, $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$. By the chosen parameters of $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, and $I_{\not\supset}$ (see Table 4 and Figure 3) it holds that (iii) $I_\subset \cup I_\cap \cup I_{\not\supset} \subseteq I_\cap$ (also $\emptyset \subseteq I_\cap$). If $\eta_\mathcal{I}(C)[i_E]$ is $\emptyset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$ then by (i)-(iii) $\eta_{\mathcal{I}}(C)[i_E] \subseteq \eta_{\mathcal{I}}(D)[i_E]$. Otherwise $\eta_{\mathcal{I}}(C)[i_E]$ is either $I_=$ or $I_\supset$. The former case happens iff $C^{\mathcal{I}} = E^{\mathcal{I}} \neq \emptyset$. However, by the ($\Rightarrow$) assumption, we have that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. So $C^{\mathcal{I}} = E^{\mathcal{I}} \neq \emptyset$ implies $\emptyset \neq E^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, which cannot happen by the "none of the above" assumption in this case. The latter case happens iff $C^{\mathcal{I}} \supset E^{\mathcal{I}} \neq \emptyset$. Again by the ($\Rightarrow$) assumption, we have that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. So $C^{\mathcal{I}} \supset E^{\mathcal{I}} \neq \emptyset$ implies $\emptyset \neq E^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, which cannot happen by the "none of the above" assumption.

It remains to argue for the dimensions of the form $i_{R,c}$, where $R \in \mathsf{N_R}$ and $c \in \Delta^{\mathcal{I}}$. By Definition 6, for every pair (R, c) with $R \in \mathsf{N_R}$ and $c \in \Delta^{\mathcal{I}}$, $\eta_{\mathcal{I}}(C)[i_{R,c}] = \eta_{\mathcal{I}}(D)[i_{R,c}] = I$. So, for every dimension of this form, $\eta_{\mathcal{I}}(C)[i_{R,c}] \subseteq \eta_{\mathcal{I}}(D)[i_{R,c}]$.

($\Leftarrow$) Now suppose $\eta_{\mathcal{I}}(C) \subseteq \eta_{\mathcal{I}}(D)$. If $\eta_{\mathcal{I}}(C) = \emptyset$ then, by Definition 6, this happens iff $C^{\mathcal{I}} = \emptyset$ and we are done since this trivially implies $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. If $\eta_{\mathcal{I}}(D) = \emptyset$ then, by the ($\Leftarrow$) assumption, $\eta_{\mathcal{I}}(C) = \emptyset$. Since, again by Definition 6, this happens iff $C^{\mathcal{I}} = \emptyset$ we are done since this trivially implies $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. We can then assume that both $\eta_{\mathcal{I}}(C) \neq \emptyset$ and $\eta_{\mathcal{I}}(D) \neq \emptyset$ hold, which means by Definition 6 that $C^{\mathcal{I}} \neq \emptyset$ and $D^{\mathcal{I}} \neq \emptyset$. By Definition 6, $\eta_{\mathcal{I}}(D)[i_D] = I_=$. By the ($\Leftarrow$) assumption, (i) $\eta_{\mathcal{I}}(C)[i_D] \subseteq I_=$. By Definition 6 $\eta_{\mathcal{I}}(C)[i_D]$ is either equal to $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, or $I_{\not\supset}$. By the chosen parameters of $I_=$, $I_\supset$, $I_\subset$, $I_\cap$, and $I_{\not\supset}$ (see Table 4 and Figure 3) the only options that are subsets of $I_=$ are $I_=$ itself and $I_\subset$. By (i), either $\eta_{\mathcal{I}}(C)[i_D] = I_=$ or $\eta_{\mathcal{I}}(C)[i_D] = I_\subset$. The former case can only happen if $C^{\mathcal{I}} = D^{\mathcal{I}}$, which means that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, as required. The latter case can only happen if $C^{\mathcal{I}} \subset D^{\mathcal{I}}$, which means that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, as required. □

Lemma 4. *Given an interpretation $\mathcal{I}$ with finite domain, let $\eta_{\mathcal{I}}$ be as in Definition 6. For all DL-Lite$^{\mathcal{H}}$ RIs $R \sqsubseteq S$, $\mathcal{I} \models R \sqsubseteq S$ iff $\eta_{\mathcal{I}} \models R \sqsubseteq S$.*

Proof. ($\Rightarrow$) Assume $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. To show that $\eta_{\mathcal{I}}(R) \subseteq \eta_{\mathcal{I}}(S)$, we show that

$$\mathsf{Head}_{\mathcal{I}}(R) \subseteq \mathsf{Head}_{\mathcal{I}}(S),$$
$$\mathsf{Tail}_{\mathcal{I}}(R) \subseteq \mathsf{Tail}_{\mathcal{I}}(S),$$
$$\mathsf{Bump}_{\mathcal{I}}(R) \subseteq \mathsf{Bump}_{\mathcal{I}}(S).$$

We start with arguing that $\mathsf{Head}_{\mathcal{I}}(R) \subseteq \mathsf{Head}_{\mathcal{I}}(S)$. We make the following case distinction.

- **Dimension i_C.** By assumption, $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. This implies $(\exists R)^{\mathcal{I}} \subseteq (\exists S)^{\mathcal{I}}$. By Definition 6, $\mathsf{Head}_{\mathcal{I}}(R)[i_C] = \eta_{\mathcal{I}}(\exists R)[i_C]$ and $\mathsf{Head}_{\mathcal{I}}(S)[i_C] = \eta_{\mathcal{I}}(\exists S)[i_C]$. By Lemma 3, $\eta_{\mathcal{I}}(\exists R)[i_C] \subseteq \eta_{\mathcal{I}}(\exists S)[i_C]$ and we are done.
- **Dimension $i_{T,c}$ and for all $e \in \Delta^{\mathcal{I}}$ we have that $(c, e) \in S^{\mathcal{I}}$ implies $(c, e) \in T^{\mathcal{I}}$.** By Definition 6, $\mathsf{Head}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{S}_\subset$. By assumption, $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. Then, for all $e \in \Delta^{\mathcal{I}}$ we have that $(c, e) \in R^{\mathcal{I}}$ implies $(c, e) \in T^{\mathcal{I}}$. By Definition 6, $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset$. So $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] = \mathsf{Head}_{\mathcal{I}}(S)[i_{T,c}]$ and thus $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Head}_{\mathcal{I}}(S)[i_{T,c}]$.
- **Dimension $i_{T,c}$ and there is $e \in \Delta^{\mathcal{I}}$ with $(c, e) \in S^{\mathcal{I}}$ but $(c, e) \notin T^{\mathcal{I}}$.** By Definition 6, $\mathsf{Head}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{S}_\subset^\neg$. Also, $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset$ or $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset^\neg$. In both cases, by the values in Table 6, we have that $\mathsf{Head}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Head}_{\mathcal{I}}(S)[i_{T,c}]$.

We now argue that $\mathsf{Tail}_{\mathcal{I}}(R) \subseteq \mathsf{Tail}_{\mathcal{I}}(S)$.

- **Dimension i_C.** By assumption, $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. This implies $(R^-)^{\mathcal{I}} \subseteq (S^-)^{\mathcal{I}}$, so $(\exists R^-)^{\mathcal{I}} \subseteq (\exists S^-)^{\mathcal{I}}$. By Definition 6, $\mathsf{Tail}_{\mathcal{I}}(R)[i_C] = \eta_{\mathcal{I}}(\exists R^-)[i_C]$ and $\mathsf{Tail}_{\mathcal{I}}(S)[i_C] = \eta_{\mathcal{I}}(\exists S^-)[i_C]$. Then, by Lemma 3, $\eta_{\mathcal{I}}(\exists R^-)[i_C] \subseteq \eta_{\mathcal{I}}(\exists S^-)[i_C]$ and we are done.
- **Dimension $i_{T,c}$ and for all $e \in \Delta^{\mathcal{I}}$ we have that $(e, c) \in S^{\mathcal{I}}$ implies $(c, e) \in T^{\mathcal{I}}$.** In this case, $\mathsf{Tail}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{S}_\subset$. By assumption, $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. Then, for all $e \in \Delta^{\mathcal{I}}$ we have that $(e, c) \in R^{\mathcal{I}}$ implies $(c, e) \in T^{\mathcal{I}}$. By Definition 6, $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset$. So $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] = \mathsf{Tail}_{\mathcal{I}}(S)[i_{T,c}]$ and thus $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Tail}_{\mathcal{I}}(S)[i_{T,c}]$.
- **Dimension $i_{T,c}$ and there is $e \in \Delta^{\mathcal{I}}$ with $(e, c) \in S^{\mathcal{I}}$ but $(c, e) \notin T^{\mathcal{I}}$.** By Definition 6, $\mathsf{Tail}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{S}_\subset^\neg$. Also, $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset$ or $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{S}_\subset^\neg$. In both cases, by the values in Table 6, we have that $\mathsf{Tail}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Tail}_{\mathcal{I}}(S)[i_{T,c}]$.

Finally, we argue that $\mathsf{Bump}_{\mathcal{I}}(R) \subseteq \mathsf{Bump}_{\mathcal{I}}(S)$.

- **Dimension i_C.** In this case, by Definition 6, $\mathsf{Bump}_{\mathcal{I}}(R)[i_C] = \mathsf{Bump}_{\mathcal{I}}(S)[i_C] = I_0$. Then, trivially, $\mathsf{Bump}_{\mathcal{I}}(R)[i_C] \subseteq \mathsf{Bump}_{\mathcal{I}}(S)[i_C]$.
- **Dimension $i_{S,c}$ and for all $e \in \Delta^{\mathcal{I}}$ we have that $(c, e) \in R^{\mathcal{I}}$ implies $(c, e) \in T^{\mathcal{I}}$.** By Definition 6, $\mathsf{Bump}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{B}_{\mathsf{R,I}}$. Also, $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{B}_{\mathsf{R,I}}$ or $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{B}_{\mathsf{R,I}}^\neg$. In both cases, by the values in Table 6, we have that $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Bump}_{\mathcal{I}}(S)[i_{T,c}]$.
- **Dimension $i_{T,c}$ and there is $e \in \Delta^{\mathcal{I}}$ with $(c, e) \in R^{\mathcal{I}}$ but $(c, e) \notin T^{\mathcal{I}}$.** By Definition 6, $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] = \mathcal{B}_{\mathsf{R,I}}^\neg$. By assumption, we have that $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. This implies that there is $e \in \Delta^{\mathcal{I}}$ with $(c, e) \in S^{\mathcal{I}}$ but $(c, e) \notin T^{\mathcal{I}}$. By Definition 6, $\mathsf{Bump}_{\mathcal{I}}(S)[i_{T,c}] = \mathcal{B}_{\mathsf{R,I}}^\neg$. So $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] = \mathsf{Bump}_{\mathcal{I}}(S)[i_{T,c}]$ and thus $\mathsf{Bump}_{\mathcal{I}}(R)[i_{T,c}] \subseteq \mathsf{Bump}_{\mathcal{I}}(S)[i_{T,c}]$.

($\Leftarrow$) Assume $\eta_{\mathcal{I}}(R) \subseteq \eta_{\mathcal{I}}(S)$. Let c be an arbitrary element of $\Delta^{\mathcal{I}}$. By Definition 6, $\mathsf{Head}_{\mathcal{I}}(S)[i_{S,c}] = \mathcal{S}_\subset$. By Definition 4 and the ($\Leftarrow$) assumption, $\mathsf{Head}_{\mathcal{I}}(R) \subseteq \mathsf{Head}_{\mathcal{I}}(S)$, so (i) $\mathsf{Head}_{\mathcal{I}}(R)[i_{S,c}] \subseteq \mathcal{S}_\subset$. By Definition 6 again, $\mathsf{Head}_{\mathcal{I}}(R)$ can be either $\mathcal{S}_\subset$ or $\mathcal{S}_\subset^\neg$. Since $\mathcal{S}_\subset^\neg \not\subseteq \mathcal{S}_\subset$ and (i) holds, we actually have that (ii) $\mathsf{Head}_{\mathcal{I}}(R)[i_{S,c}] = \mathcal{S}_\subset$ and, by Definition 6, for all $e \in \Delta^{\mathcal{I}}$, $(c, e) \in R^{\mathcal{I}}$ implies $(c, e) \in S^{\mathcal{I}}$. Since c was an arbitrary element of $\Delta^{\mathcal{I}}$, (iii) this holds for all such elements. Thus, by (ii)-(iii), $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$. □

Theorem 3. *There exists a suitable s_Ω such that every satisfiable DL-Lite$^{\mathcal{H}}$ KB $\mathcal{K}$ has a box interpretation $\eta_{\mathcal{I}_\mathcal{K}}$ that is a KB faithful model and box consistent.*

Proof. Given a satisfiable DL-Lite$^{\mathcal{H}}$ KB $\mathcal{K}$, let $\mathcal{I}_{\mathcal{K}}$ be the canonical model for $\mathcal{K}$, as in Definition 1. By Theorem 1, for all DL-Lite$^{\mathcal{H}}$ axioms α, we have that $\mathcal{K} \models \alpha$ iff $\mathcal{I}_{\mathcal{K}} \models \alpha$. By Lemmas 1, 2, 3, and 4, $\mathcal{I}_{\mathcal{K}} \models \alpha$ iff $\eta_{\mathcal{I}_{\mathcal{K}}} \models \alpha$. Then, $\mathcal{K} \models \alpha$ iff $\eta_{\mathcal{I}_{\mathcal{K}}} \models \alpha$, which corresponds to the notion of strong KB faithfulness (Lütfü Özçep, Leemhuis, and Wolter, 2020; Bourgaux et al., 2024). Note that $\Omega[i_C] = (-4, 4)$ (Table 6) implies $s_\Omega = 4$ in our construction. □

Corollary 1. *Let $\mathcal{K}$ be a DL-Lite$^{\mathcal{H}}$ KB with an empty ABox. Then for any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that $d_{\min} = |\mathsf{N_C}| + 3|\mathsf{N_R}|$ and η is a TBox faithful model of $\mathcal{K}$.*

Proof. By Claim 9, any DL-Lite$^{\mathcal{H}}$ KB with an empty ABox is satisfiable. Then, Corollary 1 follows from the proof of Theorem 3. Specifically, if $\mathcal{K}$'s ABox is empty, then the translation of a model $\mathcal{I}$ of $\mathcal{K}$ to $\eta_{\mathcal{I}}$ (see Definition 6) reduces to (1) one dimension i_C for any concept $C \in \mathsf{N_C^{\exists}}$ and (2) one dimension $i_{R,c}$ for any pair of roles $R \in \mathsf{N_R}$ and elements $c \in \Delta^{\mathcal{I}}$. Thus, the constructed embedding has $|\mathsf{N_C^{\exists}}| + |\mathsf{N_R}||\Delta^{\mathcal{I}}|$ dimensions. This dimensionality bound depends on $|\Delta^{\mathcal{I}}|$. Thus, translating different models of $\mathcal{K}$ leads to box interpretations with different dimensionalities. This means that the selected model of $\mathcal{K}$ influences the dimensionality bounds for weak faithfulness, as we show next.

For any $\mathcal{I} \models \mathcal{K}$, (i) by Theorem 6 it holds that the constructed $\eta_{\mathcal{I}}$ is box consistent and (ii) by Lemmas 1-4, it holds that $\eta_{\mathcal{I}}$ is KB-entailed for $\mathcal{K}$. By Points (i) and (ii) and Proposition 1 it holds that any $\eta_{\mathcal{I}}$ with $\mathcal{I} \models \mathcal{K}$ is a (weakly) TBox faithful model of $\mathcal{K}$. In particular, we can choose a model $\mathcal{I}$ with the smallest domain $|\Delta^{\mathcal{I}}| = 1$, which results in $\eta_{\mathcal{I}}$ having the following number of dimensions:

$$\begin{aligned}|\mathsf{N_C^{\exists}}| + |\mathsf{N_R}||\Delta^{\mathcal{I}}| &= \\ |\mathsf{N_C}| + 2|\mathsf{N_R}| + |\mathsf{N_R}| &= \\ |\mathsf{N_C}| + 3|\mathsf{N_R}|&\end{aligned}$$

One can extend $\eta_{\mathcal{I}_{\mathcal{K}}}$ to a box interpretation $\eta'_{\mathcal{I}_{\mathcal{K}}}$ with more dimensions while keeping $\eta'_{\mathcal{I}_{\mathcal{K}}}$ a weakly TBox faithful model of $\mathcal{K}$. This can be done, for instance, by copying an arbitrary dimension of $\eta_{\mathcal{I}_{\mathcal{K}}}$ multiple times. This keeps the extended $\eta'_{\mathcal{I}_{\mathcal{K}}}$ a weakly TBox faithful model of $\mathcal{K}$, while increasing its dimensionality arbitrarily. Thus, we have shown that for any $d \geq |\mathsf{N_C}| + 3|\mathsf{N_R}|$ there is a η, s.t., η is a weakly TBox faithful model of $\mathcal{K}$. □

Corollary 2. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB. Then for any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that: $d_{\min} = |\mathsf{N_C}| + |\mathsf{N_R}|(2 + |\mathsf{N_I}| + 2|\mathsf{N_R}|)$ and η is a KB faithful model of $\mathcal{K}$.*

Proof. Analogous to the proof of Corollary 1, Corollary 2 essentially follows from the proof of Theorem 3. Specifically, the translation of a model $\mathcal{I}$ of $\mathcal{K}$ to $\eta_{\mathcal{I}}$ (see Definition 6) reduces to (1) one dimension i_C for any concept $C \in \mathsf{N_C^{\exists}}$ and (2) one dimension $i_{R,c}$ for any pair of roles $R \in \mathsf{N_R}$ and elements $c \in \Delta^{\mathcal{I}}$. Thus, the constructed embedding has $|\mathsf{N_C^{\exists}}| + |\mathsf{N_R}||\Delta^{\mathcal{I}}|$ dimensions. This dimensionality bound depends again on $|\Delta^{\mathcal{I}}|$. Thus, translating different models of $\mathcal{K}$ leads to box interpretations with different dimensionalities. This means, in particular, that the selected model of $\mathcal{K}$ influences the dimensionality bounds for faithfulness, as we show next.

For any $\mathcal{I} \models \mathcal{K}$, (i) by Theorem 6 it holds that the constructed $\eta_{\mathcal{I}}$ is box consistent and (ii) by Lemmas 1-4, it holds that $\eta_{\mathcal{I}} \models \mathcal{K}$. By Points (i) and (ii) and Proposition 2 it holds that any $\eta_{\mathcal{I}}$ with $\mathcal{I} \models \mathcal{K}$ is weakly KB faithful. In particular, we know by Corollary 3 that $\mathcal{K}$ has a model $\mathcal{I}$ with $|\Delta^{\mathcal{I}}| = |\mathsf{N_I}| + 2|\mathsf{N_R}|$, which results in $\eta_{\mathcal{I}}$ having the following number of dimensions:

$$\begin{aligned}|\mathsf{N_C^{\exists}}| + |\mathsf{N_R}||\Delta^{\mathcal{I}}| &= \\ |\mathsf{N_C}| + 2|\mathsf{N_R}| + |\mathsf{N_R}|(|\mathsf{N_I}| + 2|\mathsf{N_R}|) &= \\ |\mathsf{N_C}| + |\mathsf{N_R}|(2 + |\mathsf{N_I}| + 2|\mathsf{N_R}|)&\end{aligned}$$

As argued in Corollary 1, one can modify $\eta_{\mathcal{I}_{\mathcal{K}}}$ so as to add more dimensions while keeping it a weakly KB faithful model of $\mathcal{K}$. Thus, we have shown that for any $d \geq |\mathsf{N_C}| + |\mathsf{N_R}|(2 + |\mathsf{N_I}| + 2|\mathsf{N_R}|)$ there is a η, s.t., η is a weakly KB faithful model of $\mathcal{K}$. □

E Convex Optimization Formulation for BoxLitE and Proofs for Section 5

Here we provide the details for the signed distance function and the convex optimization problem in Equation (8).

E.1 Signed distance function

The notation here is as described in Section 5.2.

Proposition 3. *The signed distance to $\mathbb{R}^n_-$ satisfies*

$$\mathsf{sdist}(y, \mathbb{R}^n_-) = \begin{cases} \|y^+\|_2 & \text{if } y \notin \mathbb{R}^n_- \\ \max_{i \in \{1,\dots,n\}} y_i & \text{if } y \in \mathbb{R}^n_-, \end{cases} \tag{5}$$

where $\mathbb{R}^n_- = \{y \in \mathbb{R}^n \mid y_i \leq 0, 1 \leq i \leq n\}$.

Proof. This proposition also follows from Example 5.1 in (Luo, Wang, and Lukens, 2018), where the expression for $\mathsf{sdist}(y, \mathbb{R}^n_+)$ is described. For the sake of self-containment, we present a complete proof here.

We consider two cases. First, suppose $y \notin \mathbb{R}^n_-$, so, by definition, $\mathsf{sdist}(y, \mathbb{R}^n_-) = \mathsf{dist}_e(y, \mathbb{R}^n_-)$. By direct computation,

$$\begin{aligned}\mathsf{dist}_e(y, \mathbb{R}^n_-) &= \inf_{x \in \mathbb{R}^n_-} \left[\sum_{i=1}^n (y_i - x_i)^2\right]^{1/2} \\ &= \left[\sum_{i=1}^n \inf_{x_i \in \mathbb{R}_-} (y_i - x_i)^2\right]^{1/2} \\ &= \left[\sum_{i=1}^n \max(y_i, 0)^2\right]^{1/2} \\ &= \|y^+\|_2.\end{aligned}$$

This takes care of the first case.

Next, suppose that $y \in \mathbb{R}^n_-$. Let $\partial\mathbb{R}^n_-$ denote the topological boundary of $\mathbb{R}^n_-$, which corresponds to the elements of $\mathbb{R}^n_-$ that have at least one zero component. As $y \in \mathbb{R}^n_-$ by assumption, the distance between y and $(\mathbb{R}^n_-)^c$ is the same as the distance between y and the boundary $\partial\mathbb{R}^n_-$, that is, $\mathsf{dist}_e(y, (\mathbb{R}^n_-)^c) = \mathsf{dist}_e(y, \partial\mathbb{R}^n_-)$.

Let $\mathcal{P}$ denote the set of all subsets of $\{1, \ldots, n\}$ except for $\emptyset$. Then, $\partial\mathbb{R}^n_-$ is the union of sets of the form $\mathbb{R}^n_{\mathcal{Q}} := \{x \in \mathbb{R}^n_- \mid x_i = 0, \forall i \in \mathcal{Q}\}$, where $\mathcal{Q} \in \mathcal{P}$. With that, $\mathsf{dist}_e(y, (\mathbb{R}^n_-)^c)$ is the minimum of the distances between y and all $\mathbb{R}^n_{\mathcal{Q}}$ for $\mathcal{Q} \in \mathcal{P}$.

Given $\mathcal{Q} \in \mathcal{P}$, we have that $\mathsf{dist}_e(y, \mathbb{R}^n_{\mathcal{Q}})$ is equal to

$$\inf_{x \in \mathbb{R}^n_{\mathcal{Q}}} \left[\sum_{i \in \mathcal{Q}} (y_i - x_i)^2 + \sum_{i \in \{1,\ldots,n\} \setminus \mathcal{Q}} (y_i - x_i)^2 \right]^{1/2}$$
$$= \left[\sum_{i \in Q} y_i^2 \right]^{1/2},$$

where the equality follows because $x_i = 0$ for $i \in Q$ and the infimum value for the second summation is zero since x_i can be taken to be equal to y_i when i does not belong to Q.

Therefore, in order to minimize $\mathsf{dist}_e(y, \mathbb{R}^n_{\mathcal{Q}})$ over all $\mathcal{Q} \in \mathcal{P}$, it is enough to consider a singleton subset $\mathcal{Q}$ corresponding to a component of y that has the smallest absolute value. Since $y \in \mathbb{R}^n_-$, the absolute value of a smallest component of y is given by $|\max_{i \in \{1,\ldots,n\}} y_i|$. Overall

$$-\mathsf{dist}_e(y, (\mathbb{R}^n_-)^c) = -\mathsf{dist}_e(y, \partial\mathbb{R}^n_-) = \max_{i \in \{1,\ldots,n\}} y_i.$$

□

As CVXPY does not implement signed distance functions natively, we implement $\mathsf{sdist}(\,\cdot\,, \mathbb{R}^n_-)$ using the support function feature provided by CVXPY. We now review some convex analysis concepts related to that.

Given a set $S \subseteq \mathbb{R}^n$, we define the *support function of* S by $\sigma_S(y) := \sup\{\langle y, z\rangle \mid z \in S\}$, where $\langle y, z\rangle$ indicates the usual Euclidean dot product between y and z.

The signed distance function to a convex set is convex, e.g., see Theorem 10.1 in Chapter 7 of (Delfour and Zolésio, 2011) or Section 3.3 of (Luo, Wang, and Lukens, 2018). In particular, since $\mathbb{R}^n_-$ is a convex cone[9], the function $\mathsf{sdist}(\cdot, \mathbb{R}^n_-)$ is convex and positively homogeneous[10]. Furthermore, $\mathsf{sdist}(\cdot, \mathbb{R}^n_-)$ is finite everywhere. Then, it follows from Corollary 13.2.2 in (Rockafellar, 1997) that $\mathsf{sdist}(\cdot, \mathbb{R}^n_-)$ can be expressed as the support function of a certain convex set. That is, there exists a convex set $C \subseteq \mathbb{R}^n$ such that $\mathsf{sdist}(y, \mathbb{R}^n_-) = \sigma_C(y)$ holds for every $y \in \mathbb{R}^n$. We implement $\mathsf{sdist}(\,\cdot\,, \mathbb{R}^n_-)$ in our code by expressing it as the support function of a certain convex set.

Let $B_n := \{y \in \mathbb{R}^n \mid \|y\|_2 \le 1\}$ denote the unit ball in $\mathbb{R}^n$ and $P_n := \{y \in \mathbb{R}^n \mid y \ge 0, y_1 + \cdots y_n \ge 1\}$. In what follows, given $y, z \in \mathbb{R}^n$, we use $y \le z$ to indicate that $y_i \le z_i$ holds for $i \in \{1, \ldots, n\}$. With that, the next proposition tell us precisely how to obtain the signed distance function to $\mathbb{R}^n_-$ as the support function of a convex set.

Proposition 4. *Let $C_n = B_n \cap P_n$. For every $y \in \mathbb{R}^n$ the following hold.*

$$\begin{aligned}\mathsf{sdist}(y, \mathbb{R}^n_-) &= \sigma_{C_n}(y) \\ &= \inf_{z \in \mathbb{R}^n} \{\|z\|_2 + \max_{i \in \{1,\ldots,n\}} (y_i - z_i) \mid y \le z\},\end{aligned}$$

Proof. The result is straightforward for $n = 1$, so henceforth we assume that $n \ge 2$. Before we proceed, we need some extra convex analysis preliminaries, for more details see (Rockafellar, 1997; Hiriart-Urruty and Lemaréchal, 1993a,b).

For a set $S \subseteq \mathbb{R}^n$ denote its indicator function by $\delta_S : \mathbb{R}^n \to \mathbb{R}^n \cup \{+\infty\}$. By definition, we have

$$\delta_S(y) := \begin{cases} 0 & \text{if } y \in S \\ +\infty & \text{if } y \notin S, \end{cases} \qquad \sigma_S(y) := \sup_{z \in S} \langle z,\, y\rangle.$$

Next, let $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ be a convex function. Its conjugate function is defined as

$$f^*(s) := \sup_{x \in \mathbb{R}^n} (\langle s,\, x\rangle - f(x)).$$

We note that $\delta_S^* = \sigma_S$, for $S \subseteq \mathbb{R}^n$.

Since $C_n = P_n \cap B_n$, we have $\delta_{B_n} + \delta_{P_n} = \delta_{C_n}$. Then, for every $y \in \mathbb{R}^n$, we have

$$\begin{aligned}\sup_{x \in C_n} \langle y,\, x\rangle &= \sup_{x \in \mathbb{R}^n} (\langle y,\, x\rangle - \delta_{B_n}(x) - \delta_{P_n}(x)) \\ &= (\delta_{B_n} + \delta_{P_n})^*(y).\end{aligned}$$

Since $n \ge 2$, the (topological) interiors of P_n and B_n intersect. For example, $(0.6, 0.6, \epsilon, \ldots, \epsilon)$ belongs to the interior of both sets for sufficiently small $\epsilon > 0$ when $n \ge 3$. For $n = 2$, it is enough to take $(0.6, 0.6)$. Under this condition, a theorem from convex analysis says that $(\delta_{B_n} + \delta_{P_n})^*$ is the exact infimal convolution between σ_{B_n} and σ_{P_n}, e.g., see Theorem 2.3.2 of Chapter X in (Hiriart-Urruty and Lemaréchal, 1993b). This means that

$$\sigma_{C_n}(y) = \sup_{x \in C_n} \langle y,\, x\rangle = \inf_{z \in \mathbb{R}^n} \{\sigma_{B_n}(z) + \sigma_{P_n}(y - z)\} \quad (11)$$

and the infimum is attained for every y. Now, for $z \in \mathbb{R}^n$ we have

$$\sigma_{B_n}(z) = \|z\|_2,$$

which follows from the Cauchy-Schwarz inequality. Defining $e := (-1, -1, \ldots, -1)$, we have for $w \in \mathbb{R}^n$

$$\begin{aligned}\sigma_{P_n}(w) &= \sup_{0 \le u, 1 \le u_1 + \cdots u_n} \langle u,\, w\rangle = \inf_{0 \le et - w, 0 \le t} -t \\ &= \begin{cases} \max_{i \in \{1,\ldots,n\}} w_i & \text{if } w \in \mathbb{R}^n_- \\ +\infty & \text{otherwise} \end{cases},\end{aligned}$$

[9] A convex cone $\mathcal{K} \subseteq \mathbb{R}^n$ is a convex set satisfying $\alpha y \in \mathcal{K}$ for all $\alpha \ge 0$ and all $y \in \mathcal{K}$.

[10] That is, $\mathsf{sdist}(\alpha y, \mathbb{R}^n_-) = \alpha\, \mathsf{sdist}(y, \mathbb{R}^n_-)$ holds, for all $\alpha \in \mathbb{R}_+, y \in \mathbb{R}^n$.

where the second equality follows from linear programming duality. The third equality holds because the constraint $0 \leq et - w$ implies that $-t \geq w_i$ for all i. So minimizing $-t$ under this constraint and the constraint that $t \geq 0$ leads to $\max_{i \in \{1,\dots,n\}} w_i$ if $w \in \mathbb{R}^n_-$. If some component of w_i is positive, then the problem is infeasible, so the infimum is $+\infty$.

Plugging the expressions for σ_{B_n} and σ_{P_n} into Equation (11) leads to

$$\sigma_{C_n}(y) = \inf_{z \in \mathbb{R}^n} \{\|z\|_2 + \max_{i \in \{1,\dots,n\}} (y_i - z_i) \mid y \leq z\}.$$

Therefore, in order to show that the proposition holds, it is enough to construct $x^* \in C_n$ and $z^* \in \mathbb{R}^n$ satisfying $y \leq z^*$ and

$$\begin{aligned} \langle y, x^* \rangle &= \|z^*\|_2 + \max_{i \in \{1,\dots,n\}} (y_i - z_i^*) \\ &= \begin{cases} \|y^+\|_2 & \text{if } y \notin \mathbb{R}^n_- \\ \max_{i \in \{1,\dots,n\}} y_i & \text{if } y \in \mathbb{R}^n_-. \end{cases} \end{aligned}$$

This can be done constructively case-by-case as follows.

(i) Suppose $y \notin \mathbb{R}^n_-$. Then, at least one component of y is positive, so $y^+ \neq 0$. Let $x^* := \frac{y^+}{\|y^+\|_2}$ and $z^* := y^+$.

Then x^* belongs to C_n, because $\|x^*\|_2 = 1$ and $x_1^* + \cdots + x_n^* = \|x^*\|_1 \geq \|x^*\|_2 = 1$. We also have $y - z^* = y^- \leq 0$, where y^- is the nonpositive part of y. Also, since at least one component of y is positive $y - z^*$ is a nonpositive vector with at least one entry equal to zero. So $\max_{i \in \{1,\dots,n\}} (y_i - z_i^*) = 0$. Overall,

$$\langle y, x^* \rangle = \|z^*\|_2 + \max_{i \in \{1,\dots,n\}} (y_i - z_i^*) = \|y^+\|_2.$$

(ii) Suppose $y \in \mathbb{R}^n_-$. Let j be an index of y associated to its largest component[11]. We let $x^* \in \mathbb{R}^n$ be such that $x_j^* := 1$ and $x_k^* := 0$ for $k \neq j$. We have $\|x^*\|_2 = 1$ and $x_1^* + \cdots x_n^* = x_j^* = 1$, so $x^* \in C$. Finally, let $z^* := 0$. With that, since $y \in \mathbb{R}^n_-$, we have $y - z^* = y \leq 0$ and

$$y_j = \langle y, x^* \rangle = \|z^*\|_2 + \max_{i \in \{1,\dots,n\}} (y_i - z_i^*) = \max_{i \in \{1,\dots,n\}} y_i.$$

□

E.2 Convex Optimization

Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a DL-Lite$^{\mathcal{H}}$ KB, let d be the embedding dimension and let $\mathbf{s_\Omega}$ and the $\epsilon > 0$ be as in Section 3. Following Section 3, a given box interpretation η associates to each individual name $a \in \mathsf{N_I}$, each concept name $C \in \mathsf{N_C}$ and each role name $R \in \mathsf{N_R}$ the following objects: two vectors $(\mathsf{pos}(a), \mathsf{bump}(a)) \in \Omega \times \Omega$, a box $\eta(C)$ and three boxes $(\mathsf{Head}(R), \mathsf{Tail}(R), \mathsf{Bump}(R))$, respectively. Each box is parameterized by two vectors in $\mathbb{R}^d$ representing lower and upper bounds.

[11]There may be multiple j's, but any will work.

We recall that we concatenate all the parameters of η into a single vector z of dimension $n := (2|\mathsf{N_I}| + 2|\mathsf{N_C}| + 6|\mathsf{N_R}|)d$. With that, let $\mathcal{C}_\mathcal{K} \subseteq \mathbb{R}^n$ be the set of z's such that the constraints defined in Section 5.1 are satisfied. That is, $z \in \mathcal{C}_\mathcal{K}$ if and only if the η corresponding to z is such that: for each TBox axiom the corresponding inequalities are satisfied; and the box consistency and universe constraints are satisfied. Also, let $f_{\boldsymbol{\lambda}} : \mathbb{R}^n \to \mathbb{R}$ be the function that maps z to the objective value described in Section 5.2 for a given choice of nonnegative hyperparameters $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3_+$.

In what follows, we recall that a set in $\mathbb{R}^n$ is said to be *polyhedral* if it can be written as the set of solutions of finitely many linear equalities/inequalities.

Theorem 4. *Let $\mathcal{K}$ be a satisfiable DL-Lite$^{\mathcal{H}}$ KB. For nonnegative $\boldsymbol{\lambda}$ the following optimization problem is convex.*

$$\min_{z \in \mathbb{R}^n} f_{\boldsymbol{\lambda}}(z), \quad \text{subject to } z \in \mathcal{C}_\mathcal{K}. \tag{8}$$

In particular, $f_{\boldsymbol{\lambda}}$ is a convex function, $\mathcal{C}_\mathcal{K}$ is a polyhedral set and the following items hold.

- *i) For d as in Corollary 1, and s_Ω as in Theorem 3, $\mathcal{C}_\mathcal{K}$ is nonempty.*
- *ii) Any embedding solution $z \in \mathcal{C}_\mathcal{K}$ corresponds to a box consistent interpretation that is TBox faithful.*
- *iii) If $\lambda_1 = 0$ and there is an optimal solution z^* such that $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ then the box interpretation corresponding to z^* is KB faithful.*
- *iv) Suppose that $\lambda_1 = \lambda_2 = \lambda_3 = 0$. For d as in Corollary 2, and s_Ω as in Theorem 3, there is an optimal solution z^* s.t. $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ holds.*

Proof. As described in Section 5.1, each inclusion in the TBox of the KB is translated into finitely many linear inequalities in terms of the parameters of the box interpretation, which are exactly the components of z. Similarly, for each individual, concept and role, the box consistency and universe constraints in Section 5.1 are translated to finitely many linear inequalities in z. Since $\mathsf{N_I}$, $\mathsf{N_C}$, $\mathsf{N_R}$, and the TBox are all finite, $\mathcal{C}_\mathcal{K}$ is the intersection of solution sets of finitely many linear inequalities. Thus, $\mathcal{C}_\mathcal{K}$ is a polyhedral set in $\mathbb{R}^n$.

Next, we move on to the convexity of $f_{\boldsymbol{\lambda}}$. First, we recall that the composition of a convex function with an affine function is still convex. Also, sums of convex functions are convex. Similarly, the maximum of convex functions is also convex. See (Boyd and Vandenberghe, 2004, Section 3.2) for a review of calculus rules for convex functions.

First, the signed distance function $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ is convex, see Theorem 10.1 in Chapter 7 of (Delfour and Zolésio, 2011) or pg.154 in (Hiriart-Urruty and Lemaréchal, 1993a). With that, the dist function is also convex, as it is the composition of a convex function with an affine function. This implies that the concept assertion loss $\mathcal{L}_{\mathsf{concept}}$ is convex as well as it is again a composition of a convex function with an affine function. Similarly, both the role assertion loss $\mathcal{L}_{\mathsf{role}}$ and the negative concept regularization $\mathcal{L}_{\mathsf{negative}}$ are convex, as they are the maximum of finitely many convex functions. Finally, the box width regularization $\mathcal{R}_{\mathsf{width}}$ is also convex, since it is the composition of a convex function (the 2-norm) with a linear function.

Overall the objective function $f_{\boldsymbol{\lambda}}$ is convex as it is obtained by taking sums and maximums of finitely many convex functions.

So far, we have shown that the problem in (8) is a convex optimization problem with polyhedral constraints. Next, we move on to the proofs of the items.

$\boxed{i)}$ Let $\tilde{\mathcal{K}}$ be the KB which coincides with $\mathcal{K}$ except for the fact that its ABox is empty. By item $i)$ of Corollary 1, for every $d \geq d_{\min}$, there exists a box interpretation η that is weakly TBox faithful. Because of the faithfulness of the embedding, η satisfies the constraints described in Section 5.1.

First we recall that the proof of Corollary 1 is done by constructing η following Definition 6 with $s_\Omega = 4$ (as in the proof of Theorem 3), see also Table 6. This allows us to invoke Theorem 6, which ensures that η satisfies (1) in Section 5.1. Furthermore, because η is a box interpretation, the inequalities in (2) and (3) (Section 5.1) must be satisfied, see Definition 3. This is because Definition 3 imposes the same restrictions on the widths of the boxes.

Finally, aggregating the parameters of η into a single vector $z \in \mathbb{R}^n$ in an appropriate order, we will have $z \in \mathcal{C}_{\mathcal{K}}$, which implies that $\mathcal{C}_{\mathcal{K}} \neq \emptyset$.

$\boxed{ii)}$ The proof of item $ii)$ follows from the definition of $\mathcal{C}_{\mathcal{K}}$. We recall that $z \in \mathcal{C}_{\mathcal{K}}$ if and only if the corresponding inequalities in Section 5.1 are satisfied. As in the proof of Theorem 6, the inequality in (1) (Section 5.1) implies box consistency. We conclude that the box interpretation corresponding to a given $z \in \mathcal{C}_{\mathcal{K}}$ is box consistent and satisfies all the TBox axioms of the underlying KB. By Proposition 1, the box interpretation associated to z is TBox faithful, which concludes the proof.

$\boxed{iii)}$ Since z^* is assumed to be an optimal solution, we have $z^* \in \mathcal{C}_{\mathcal{K}}$. By item $ii)$, the box interpretation corresponding to z^* satisfies all the TBox axioms and must be box consistent. Because $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ holds, λ_1 is zero and the regularization terms associated to λ_2, λ_3 are nonnegative, we conclude that the first max term of the objective function term is nonpositive. In particular, all the $\mathcal{L}_{\mathsf{concept}}$ and $\mathcal{L}_{\mathsf{role}}$ terms inside the max must be nonpositive. By the definition of the signed distance function $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_{-})$, this implies that the corresponding ABox axioms are satisfied as well. By Proposition 2, the box interpretation associated with z^* is KB faithful.

$\boxed{iv)}$ By item $i)$ of Corollary 2, there exists a weakly KB faithful box interpretation η of $\mathcal{K}$ for any $d \geq d_{\min}$. As in the proof of item $i)$, η satisfies the constraints described in Section 5.1 for $s_\Omega = 4$, which comes as consequence of Theorem 6, the proof of Corollary 2 and the definition of box interpretation in Definition 3.

Aggregating the parameters of the box interpretation η into a single vector $z^* \in \mathbb{R}^n$ in an appropriate order, we have $z^* \in \mathcal{C}_{\mathcal{K}}$. Because $\eta \models \mathcal{K}$, all the ABox axioms are satisfied, so the the $\mathcal{L}_{\mathsf{concept}}$ and $\mathcal{L}_{\mathsf{role}}$ terms in in first max term of the objective function must be nonpositive. As a consequence, if the regularization parameters $\lambda_1, \lambda_2, \lambda_3$ in the objective function term are zero as well, then $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ holds. □

Second-order cone representability. For those familiar with conic optimization, the proof of Theorem 5 is routine and can be summarized as follows. We employ the notion of the *epigraph* of a function $f : \mathbb{R}^s \to \mathbb{R} \cup \{+\infty\}$, defined as the set $\{(x, t) \in \mathbb{R}^s \times \mathbb{R} \mid f(x) \leq t\}$. With that, the epigraphs of the functions corresponding to each of the terms appearing in $f_{\boldsymbol{\lambda}}$ are *second-order cone representable* (SOCr (Lobo et al., 1998), also called CQr in (Ben-Tal and Nemirovski, 2001, Lecture 3)). Furthermore, second-order cone representability is preserved by adding functions, composition with affine functions and taking a maximum of finitely many SOCr functions, see (Ben-Tal and Nemirovski, 2001, Section 3.3). Overall, applying appropriate calculus rules, we see that the epigraph of $f_{\boldsymbol{\lambda}}$ is SOCr, which implies in particular that the problem in (8) has a SOCP formulation. For the sake of self-containment, we present a detailed proof.

Theorem 5. *For nonnegative* $\boldsymbol{\lambda}$*, the problem in* (8) *can be reformulated as an equivalent SOCP.*

"Equivalent" in the statement of Theorem 5 means that the optimal value and optimal solutions from the former can be recovered from the optimal value and optimal solutions to the latter and vice-versa.

Proof. In its most general form, a SOCP can be written as

$$\begin{aligned} \min_{y \in \mathbb{R}^s} \quad & c^T y \\ \text{subject to} \quad & y \in \mathcal{P} \\ & b_i + A_i y \in \mathcal{K}_2^{n_i}, \qquad i = 1 \ldots, m, \end{aligned} \tag{12}$$

where each $\mathcal{K}_2^{n_i} := \{(t, z) \in \mathbb{R} \times \mathbb{R}^{n_i - 1} \mid t \geq \|z\|_2\}$ is the second-order cone in $\mathbb{R}^{n_i}$, $\mathcal{P} \subseteq \mathbb{R}^s$ is a polyhedral set described via finitely many linear equalities/inequalities[12], the b_i's are vectors, A_i's are linear maps, $c \in \mathbb{R}^s$ is a fixed vector and $c^T y$ indicates the Euclidean inner product between c and y so that $c^T y = \sum_{i=1}^s c_i y_i$ holds.

Comparing the SOCP in (12) with the problem in (8) and recalling that $\mathcal{C}_{\mathcal{K}}$ is polyhedral, we see that (8) is not a SOCP only because its objective function is nonlinear. Here, we will use the common optimization trick of "dropping the objective to the constraints", e.g., see (Lobo et al., 1998, Section 2.5) or (Ben-Tal and Nemirovski, 2001, Chapter 3). The idea is as follows, if $g : \mathbb{R}^s \to \mathbb{R}$ is a real function and $S \subseteq \mathbb{R}^s$,

[12]Depending on the reference, the "standard form" of SOCPs may not include linear inequalities directly, but this can be bypassed easily since a linear inequality of the form "$a^T y \geq d$" is equivalent to the SOC constraint "$a^T y - d \in \mathcal{K}_1^2$".

then the problem "$\min_y g(y)$ subject to $y \in S$" is equivalent to "$\min_{t,y} t$ subject to $g(y) \leq t, y \in S$". Similarly, if there were several functions g_j we would have that the problem

$$\min_y \sum_{j=1}^{\ell} g_j(y) \text{ subject to } y \in S$$

is equivalent to

$$\min_{t_1,\ldots,t_\ell,y} \sum_{i=1}^{\ell} t_j \quad \text{subject to} \quad g_j(y) \leq t_j \ (j = 1, \ldots, \ell),\, y \in S \tag{13}$$

Here, the constraints "$g_j(y) \leq t_j$" simply mean that (y, t_j) belongs to the epigraph of g_j. In particular, if S and the epigraphs of g_j can be represented via finitely many linear equalities/inequalities and SOC constraints, then the problem in (13) can reformulated as a SOCP since its objective function is linear. Informally, we say that a function is SOCr if its epigraph can be represented via finitely many equalities/inequalities and SOC constraints (adding auxiliary variables if necessary).

This discussion provides a blueprint for proving that the problem in (8) can be reformulated as a SOCP. The objective function $f_{\boldsymbol{\lambda}}$ is a sum of four terms: the loss terms associated to concept and role assertions which are aggregated with a max, the negative sampling component, and two terms for width regularization. For each term that appear, we add one auxiliary variable t_j, we "drop the objective function terms to the constraints" and we argue that each resulting constraint can be written in terms of linear equalities/inequalities and SOC constraints, i.e., the epigraphs of the functions are SOCr. Naturally, if g is SOCr then the composition of g with an affine function is SOCr, e.g., see (Ben-Tal and Nemirovski, 2001, Remark 3.3.1)[13].

The building blocks for the loss term and regularization terms in $f_{\boldsymbol{\lambda}}$ boil down to two functions: the signed distance function $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ and the 2-norm function $\|\cdot\|_2$. All the four terms in $f_{\boldsymbol{\lambda}}$ are obtained by composing copies of $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ and $\|\cdot\|_2$ with affine functions and either adding them together or taking maximums. In view of our discussion so far, it is enough to establish that the epigraphs of these two functions can be written in terms of linear equalities/inequalities and SOC constraints. In other words, we need to show that both functions are SOCr.

First, we will show that $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ is SOCr. For that, we define the auxiliary function $\psi : \mathbb{R}^{2d} \times \mathbb{R}^{2d} \to \mathbb{R} \cup \{+\infty\}$ such that

$$\psi(y, z) \coloneqq \begin{cases} \|z\|_2 + \max_{i \in \{1,\ldots,2d\}}(y_i - z_i) & \text{if } y \leq z \\ +\infty & \text{otherwise.} \end{cases}$$

[13] It is enough to observe that if h is an affine function of the form $h(x) = Bx + d$, for B a linear map and d is a vector, we can obtain a SOC representation of the epigraph of $g \circ h$ by adding the linear constraint $y = Bx + d$ to a SOC representation of $\{(y, t) \mid g(y) \leq t\}$ (the epigraph of g).

The function $\psi(y, z)$ is SOCr because $\psi(y, z) \leq t$ holds for $t \in \mathbb{R}$ if and only if there exists $t_1, t_2 \in \mathbb{R}$ such that

$$\begin{aligned} t_1 + t_2 &\leq t, \\ (t_1, \|z\|_2) &\in \mathcal{K}_2^{2d+1}, \\ y_i - z_i &\leq t_2, \quad \forall i \in \{1, \ldots, 2d\} \\ y &\leq z. \end{aligned}$$

Furthermore, Proposition 4 implies that $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ is the partial minimization of ψ with respect to the second argument, i.e.,

$$\mathsf{sdist}(y, \mathbb{R}^{2d}_-) = \inf_{z \in \mathbb{R}^{2d}} \psi(y, z)$$

holds for every $y \in \mathbb{R}^{2d}$. Furthermore, the proof of Proposition 4 shows that for every y, there exists z_y such that $\mathsf{sdist}(y, \mathbb{R}^{2d}_-) = \psi(y, z_y)$ holds, i.e., the infimum is achieved for every y. The property of being SOCr is preserved by partial minimization assuming that for every y the infimum is achieved, e.g., see Section 3.3 in (Ben-Tal and Nemirovski, 2001). Therefore, $\mathsf{sdist}(\cdot, \mathbb{R}^{2d}_-)$ is SOCr as well.

Next, the function $\|\cdot\|_2$ is also SOCr, since $\|y\|_2 \leq t$ holds for $y \in \mathbb{R}^d, t \in \mathbb{R}$ if and only if $(t, y) \in \mathcal{K}_2^d$.

In conclusion, all the functions used to build the objective function $f_{\boldsymbol{\lambda}}$ are SOCr, so overall, the problem in (8) has a SOCP formulation. □

F Size of the final optimization problem

As mentioned in Section 6, a problem modelled through CVXPY is first compiled and then sent to a solver such as MOSEK. Here we report in Table 8 the number of variables and constraints of the final optimization problem that CVXPY outputs to MOSEK.

Dataset	#Variables	#Constraints
F_v1	461k	300k
F_v2	725k	475k
F_v3	1312k	863k
F_v4	3584k	2367k

Table 8: Problem sizes of the final optimization problem solved by MOSEK split by dataset.

G Experimental Details

This section describes our experimental setup, created benchmark datastes, and evaluation protocol in detail. In particular, Section G.1 contains details about the implementation of BoxLitE. Furthermore, Section G.2 discusses the properties of the created benchmark datasets F_v1-4. Next, Section G.3, describes our experimental setup, including a detailed description of the learning setup, used hardware, and selected hyperparameters. Continuing from that, Section G.4 discusses the evaluation protocol and metrics in detail. Moreover, Section G.5 shows the runtime of BoxLitE and the SGD solvers over the various Family dataset subsets. Finally, Section G.6 details the hyperparameters used for the SGD models we have shown in our experiments.

G.1 Implementation & Reproducibility

We implemented the BoxLitE (convex) second-order cone optimization problem in Python 3.12 using CVXPY for formulating the problem and MOSEK for optimizing it. The seed used for the SGD models is 6934.

G.2 Details on F_v1-4

This section contains details about the created benchmark datasets F_v1-4 of Section 6. In particular, we have created a set of datasets (F_v1-4) of varying sizes from the family dataset (Imenes, Guimarães, and Ozaki, 2023). We derived these datasets by sampling approximately $k \in \{300, 500, 1000, 3000\}$ assertions of the family dataset's ABox with forest fire sampling (Leskovec, Kleinberg, and Faloutsos, 2005), a popular sampling technique for large graphs, using a forward burning probability (pf $= 0.7$) and a backward burning probability (bf $= 0$). Furthermore, since the family dataset solely provides role assertions in its ABox, we selected any TBox inclusion of the family dataset that includes solely roles and extended the TBox by the disjointness axioms $\exists \mathsf{hasFather}^- \sqsubseteq \neg\exists \mathsf{hasMother}^-$. Figure 1 lists the TBox of the created datasets F_v1-4.

Next, we created a set of inferred assertions by (i) adding any assertion that logically follows from each dataset and (ii) removing any assertion that occurs in the ABox. We randomly split this set of inferred assertions into a validation set (20%), used for model selection, and a test set (80%), used for evaluating the performance of the selected model.

G.3 Experimental Setup

Training Setup. During the training phase, we optimized the loss described in 5.1 using MOSEK on the train set. After retrieving an embedding solution from MOSEK, we evaluated its performance on the validation set, which we used for selecting the best embedding solution (see Section G.4 for more details on the evaluation protocol). We now discuss BoxLitE's hyperparameter optimization.

Hyperparameter Optimization. We set $s_\Omega = 1$ and $\epsilon = 10^{-2}$ for all of our experiments. For the benchmark results on datasets F_v1-4, we set $d = 32$ and tuned the hyperparameters within the following ranges: $\lambda_1, \lambda_2, \lambda_3 \in \{0, 0.001, 0.003, 0.1, 0.3, 1, 3\}$. We list the best found hyperparameters for each of the datasets in Table 9.

Dataset	λ_1	λ_2	λ_3
F_v1	.001	3	.001
F_v2	.001	1	.003
F_v3	.003	0.1	.003
F_v4	.001	3	.003

Table 9: Best found hyperparameters for BoxLitE.

G.4 Evaluation Protocol

We have evaluated BoxLitE, by following the standard evaluation setting for KB completion as described by (Abboud et al., 2020; Pavlović and Sallinger, 2023b; Xiong et al., 2022).

Dataset	λ_1	λ_2	λ_3
F_v1	.0	3	.001
F_v2	.0	3	.03
F_v3	.0	.3	.0
F_v4	.0	.3	0

Table 10: Best found hyperparameters for BoxLitE1.

Dataset	λ_1	λ_2	λ_3
F_v1	.003	.0	.0
F_v2	.01	.0	1
F_v3	.01	.0	.003
F_v4	.3	.0	.003

Table 11: Best found hyperparameters for BoxLitE2.

Dataset	λ_1	λ_2	λ_3
F_v1	.003	.1	.0
F_v2	.003	1	.0
F_v3	.003	0.1	.0
F_v4	.003	.1	.0

Table 12: Best found hyperparameters for BoxLitE3.

In particular, this includes measuring the ranking quality of each role assertion $R(a, b)$ in the test set over any possible individual in the first position of the assertion, i.e., $R(a', b)$ for all $a' \in \mathsf{N_I}$, and the second position of the assertion, i.e., $R(a, b')$ for all $b' \in \mathsf{N_I}$. Furthermore, we used the standard metrics for KB completion, namely, the mean reciprocal rank (MRR) and hits at k (H@k).

As typically done in the literature (Bordes et al., 2013; Sun et al., 2019; Abboud et al., 2020; Pavlović and Sallinger, 2023b; Xiong et al., 2022), we presented the filtered versions of these merics introduced by (Bordes et al., 2013). This means in particular that for hyperparameter tuning, we evaluated each of the found BoxLitE embedding solutions on the validation set and excluded any assertion from the ranking that occurs in the train or validation set (apart from the validation assertion whose score shall be computed). We selected those embedding solutions that reached the highest scores on the validation set. We followed the filtered setting (Bordes et al., 2013) also during the final evaluation, i.e., we evaluated the selected embedding solutions on the test set and excluded any assertion from the ranking that occurs in the train, validation, or test set (apart from the test assertion whose score shall be computed).

The intuition of the filtered setting is that we exclude assertions from the ranking that are during the current evaluation stage known to be true, as assigning a high score to these assertions does not indicate a wrong inference. Specifically, during hyperparameter tuning on the validation set, the train and validation assertions are known to be true and thus need to be excluded; while during the final evaluation on the test set, the train, validation, and test assertions are known to be true and thus need to be excluded from the ranking. Finally, we briefly review the definition of H@k and the MRR: H@k reflects the proportion of true assertions within the predicted assertions whose rank is at most k, whereas the MRR represents their average of inverse ranks ($1/rank$).

G.5 Running Time

For each choice of the hyperparameters $\lambda_1, \lambda_2, \lambda_3$, the evaluation of BoxLitE takes less than a minute for F_v1 and F_v2. It takes 2 minutes and 38 seconds for F_v3 and more than 20 minutes for F_v4. Improving our evaluation procedure would contribute for scalability.

Dataset	BoxE	RotatE	ComplEx
F_v1	8m38s	5m56s	7m19s
F_v2	8m40s	4m39s	8m37s
F_v3	44m44s	4m24s	8m30s
F_v4	34m37s	13m14s	16m27s

Table 13: Training time split by dataset for SGD methods.

Regarding SGD methods, we provide in Table 13 the training time for the SGD methods with 100 trials and 500 epochs for each method. We used the same dimensionality 32 as in BoxLitE (see G.3).

G.6 SGD setup

The SGD models shown in our experiments were trained using PyKEEN's implementation of BoxE, RotatE, and ComplEx. We trained each model on the different Family subsets for 100 trials over 500 epochs. The optimizer used is Adagrad, together with early stopping. The frequency, patience, and relative delta of the stopper are 10, 10, and 0.01, respectively. Lastly, the embedding dimension, as in BoxLitE, is set to 32. The remaining parameters are set to the defaults of the KGE implementations in PykEEN (Ali et al., 2021).