Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning
Reasoning, Learning & Decision-Making Track

Knowledge Base Embeddings: Semantics and Theoretical Properties

Camille Bourgaux[1] , **Ricardo Guimarães**[2] , **Raoul Koudijs**[2] , **Victor Lacerda**[2] , **Ana Ozaki**[2,3]
[1] DI ENS, ENS, CNRS, PSL University & Inria, Paris, France
[2] University of Bergen
[3] University of Oslo

Abstract

Research on knowledge graph embeddings has recently evolved into *knowledge base* embeddings, where the goal is not only to map facts into vector spaces but also constrain the models so that they take into account the relevant conceptual knowledge available. This paper examines recent methods that have been proposed to embed knowledge bases in description logic into vector spaces through the lens of their geometric-based semantics. We identify several relevant theoretical properties, which we draw from the literature and sometimes generalize or unify. We then investigate how concrete embedding methods fit in this theoretical framework.

1 Introduction

Knowledge graph (KG) embeddings allow for a continuous representation of KGs in vector spaces, which can be used for link prediction and related tasks. Recent works have expanded this idea to *knowledge base* (KB) embeddings, which take into account not only facts but also conceptual knowledge, expressed as a *TBox* (Gutiérrez-Basulto and Schockaert, 2018; Kulmanov et al., 2019; Özçep, Leemhuis, and Wolter, 2020; Abboud et al., 2020; Mondal, Bhatia, and Mutharaju, 2021; Peng et al., 2022; Xiong et al., 2022; Pavlovic and Sallinger, 2023; Jackermeier, Chen, and Horrocks, 2024). Which theoretical properties are interesting for KB embeddings? Which embedding methods have these properties? How expressive is the ontology language considered? These are some of the relevant questions to better understand how embedding methods work and which properties they offer.

One of the challenges to study KB embeddings in a uniform way is that the methods differ not only in how they are defined but also in the ontology language and in the properties the authors consider. We focus on KBs that can be expressed in *description logic* (DL), and on *region-based embedding methods*, which usually come with a geometric-based semantics. Regarding the properties, a basic goal is to determine whether there is some kind of correspondence between classical models based on interpretations and geometric-based models created by the embedding methods. A simple kind of correspondence is whether the existence of a (geometric-based) model within the embedding method implies the KB is satisfiable, and vice-versa, whether the existence of a classical interpretation that satisfies a given KB implies the existence

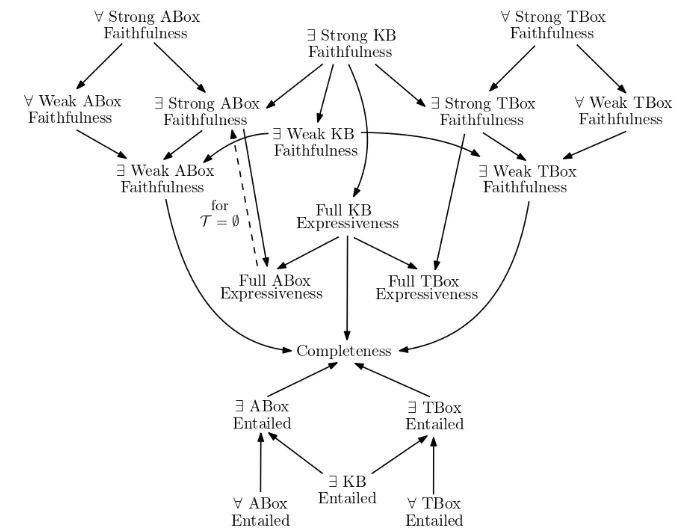

Figure 1: Relationships between the properties we consider (except for soundness which is incomparable and KB properties expressible by combining TBox and ABox properties). An arrow from property X to property Y indicates that property X implies property Y. A dashed line indicates that the implication holds when the TBox $\mathcal{T}$ is empty. The symbol ∀ defines a guarantee, while ∃ just posits ability.

of a model within the embedding method. The former property is known as *soundness* (see, e.g., Xiong et al. (2022)) and we call the latter *completeness*. Such correspondence does not require, for example, that (i) axioms entailed by a given KB hold in the geometric-based model, or conversely, that (ii) axioms that hold in the geometric-based model are a consequence of the KB or (iii) are at least consistent with the KB. These properties strengthen the notion of completeness. We call Property (i) *entailment closure*, while Properties (ii) and (iii) correspond, respectively, to the notions of *strong* and *weak faithfulness* by Özçep, Leemhuis, and Wolter (2020). We study two variants for (i)-(iii): one only requires the *ability* of an embedding method to produce a geometric-based model with the desired property, that is, whether such a model *exists*; and one where the property should hold as a *guarantee*, that is, in addition to ability, *every* model should have the property. In the KG literature, *full expressiveness* (Kazemi

Name	Syntax	Semantics
Top	$\top$	$\Delta^{\mathcal{I}}$
Bottom	$\bot$	$\emptyset$
Nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
Negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
Conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
Q. exist. res.	$\exists R.C$	$\{d \mid (d,e) \in R^{\mathcal{I}}, e \in C^{\mathcal{I}}\}$
Q. univ. res.	$\forall R.C$	$\{d \mid (d,e) \in R^{\mathcal{I}} \Rightarrow e \in C^{\mathcal{I}}\}$
Inverse	r^-	$\{(e,d) \mid (d,e) \in r^{\mathcal{I}}\}$
Negation	$\neg R$	$(\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}) \setminus R^{\mathcal{I}}$
Composition	$R \circ S$	$\{(d,e) \mid (d,d') \in R^{\mathcal{I}}, (d',e) \in S^{\mathcal{I}}\}$

Table 1: Syntax and semantics of common DL constructors: $a \in \mathsf{N_I}$, $r \in \mathsf{N_R}$, C, D are (complex) concepts, and R, S (complex) roles.

Name	Syntax	Semantics
Concept inclusion	$C_1 \sqsubseteq C_2$	$C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$
Role inclusion	$R \sqsubseteq S$	$R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$
Concept assertion	$A(a)$	$a^{\mathcal{I}} \in A^{\mathcal{I}}$
Role assertion	$r(a,b)$	$(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$

Table 2: Syntax and semantics of common TBox and ABox axioms: $A \in \mathsf{N_C}$, $r \in \mathsf{N_R}$, $a, b \in \mathsf{N_I}$, C_1, C_2 denote (complex) concepts and R, S (complex) roles.

Name	Rule form	DL
Symmetry	$\forall \mathbf{x}\, r(x,y) \rightarrow r(y,x)$	$r \sqsubseteq r^-$
Inversion	$\forall \mathbf{x}\, r(x,y) \leftrightarrow s(y,x)$	$r \equiv s^-$
Hierarchy	$\forall \mathbf{x}\, r(x,y) \rightarrow s(x,y)$	$r \sqsubseteq s$
Intersection	$\forall \mathbf{x}\, r(x,y) \wedge s(x,y) \rightarrow t(x,y)$	$r \sqcap s \sqsubseteq t$
Composition	$\forall \mathbf{x}\, r(x,y) \wedge s(y,z) \rightarrow t(x,z)$	$r \circ s \sqsubseteq t$
Mut. exclusion	$\forall \mathbf{x}\, r(x,y) \wedge s(x,y) \rightarrow \bot$	$r \sqsubseteq \neg s$
Asymmetry	$\forall \mathbf{x}\, r(x,y) \wedge r(y,x) \rightarrow \bot$	$r \sqsubseteq \neg r^-$

Table 3: Common patterns and their DL counterparts, where r, s, t are distinct roles in $\mathsf{N_R}$ and $\forall \mathbf{x}$ is a shorthand for $\forall xy$ or $\forall xyz$.

and Poole, 2018) means that, given any assignment of truth values for facts, there is an embedding model that separates true facts from false ones. We generalize this notion to include ontology languages. We study and formalize these different properties, proposing a theoretical framework for better understanding KB embeddings behaviour. Figure 1 illustrates the relationships between the properties. We also study recent KB embedding methods and investigate how they fit in the theoretical framework. Our study reveals that for many embedding methods, in particular those with an implementation, the theoretical properties stated in the literature do not hold or cannot be combined (e.g., an embedding method can be fully expressive and able to capture some patterns but not within the same model).

We provide basic definitions in Section 2 and present recent region-based embedding methods and their semantics in Section 3. In Section 4, we introduce embedding method properties and show how they relate. We also show that if the ontology language is *finite* (that is, only finitely many axioms exist in the language), which is a common assumption for KB embedding methods, then multiple properties become equivalent. In Section 5, we investigate whether the embedding methods of Section 3 fit into the theoretical framework of Section 4. We conclude in Section 6. Omitted proofs are available in (Bourgaux et al., 2024).

2 Basic Definitions

This section recalls the basics of DL syntax and semantics and the basics of KB embeddings into vector spaces.

2.1 Description Logic Knowledge Bases

Syntax Let $\mathsf{N_C}$, $\mathsf{N_R}$, and $\mathsf{N_I}$ be pairwise disjoint finite sets of *concept names*, *role names*, and *individual names* or *entities*, respectively. These sets are usually countably infinite in the DL literature (Baader et al., 2017) but often assumed to be finite in the KG and KB embedding literature (Abboud et al., 2020; Xiong et al., 2022). An *ABox* $\mathcal{A}$ is a finite set of concept and role assertions of the form $A(a)$ or $r(a,b)$ respectively, where $A \in \mathsf{N_C}$, $r \in \mathsf{N_R}$ and $a, b \in \mathsf{N_I}$. A *TBox* $\mathcal{T}$ is a finite set of axioms whose form depends on the specific DL language. Syntax of common DL constructors and TBox axioms is given in Tables 1 and 2. In the KG embedding literature, *(inference) patterns* are often considered. Table 3 presents these patterns and their DL translation. We say that a DL language $\mathcal{L}$ is *finite* if there are finitely many axioms expressible in $\mathcal{L}$. A DL *knowledge base* $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is the union of a TBox and an ABox.

Semantics The semantics of DL KBs is given by interpretations. An *interpretation* $\mathcal{I}$ is a pair $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where the *interpretation domain* $\Delta^{\mathcal{I}}$ is a non-empty set and $\cdot^{\mathcal{I}}$ is a function that maps each $a \in \mathsf{N_I}$ to some $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each $A \in \mathsf{N_C}$ to some $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and each $r \in \mathsf{N_R}$ to some $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The function $\cdot^{\mathcal{I}}$ is extended to complex concept and roles as explained in Table 1 and the satisfaction of TBox axioms and ABox assertions is defined by Table 2. An interpretation $\mathcal{I}$ is a *model* of an ABox $\mathcal{A}$ ($\mathcal{I} \models \mathcal{A}$) if it satisfies every assertion in $\mathcal{A}$; it is a model of a TBox $\mathcal{T}$ ($\mathcal{I} \models \mathcal{T}$) if it satisfies every axiom in $\mathcal{T}$; and it is a model of a KB $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ if $\mathcal{I} \models \mathcal{T}$ and $\mathcal{I} \models \mathcal{A}$. A KB $\mathcal{K}$ is *satisfiable* (or *consistent*) if it has a model. An axiom α (being an ABox assertion or a TBox axiom) is *consistent with a KB* $\mathcal{K}$ if $\mathcal{K} \cup \{\alpha\}$ is satisfiable, and is *entailed* by $\mathcal{K}$, written $\mathcal{K} \models \alpha$, if $\mathcal{I} \models \alpha$ for every model $\mathcal{I}$ of $\mathcal{K}$. The *deductive closure* of a KB $\mathcal{K}$ is the (possibly infinite) set of all axioms entailed by $\mathcal{K}$.

ABoxes as TBoxes Some KB embedding methods operate on the TBox only and encode the ABox into the TBox using nominals. Specifically, $A(a)$ is represented by $\{a\} \sqsubseteq A$ and $r(a,b)$ by $\{a\} \sqsubseteq \exists r.\{b\}$. When discussing the embedding methods properties, we still regard these axioms as assertions.

2.2 Embedding KBs Into Vector Spaces

Vector spaces, regions and transformations The aim of KG or KB embedding is to learn a low-dimensional representation of the KG or KB components into some *vector space(s)*. The d-dimensional vector space $\mathbb{R}^d$ is an Euclidean space whose elements are of the form $\vec{v} = (v_1, \ldots, v_d)$ and may be added together (using $+$) or multiplied by scalars (using .). We use $\vec{u} - \vec{v}$ as a shorthand for $\vec{u} + (-1).\vec{v}$ and say that $\vec{u} \leq \vec{v}$ if $u_i \leq v_i$ for every $1 \leq i \leq d$. The *distance* between $\vec{u}, \vec{v} \in \mathbb{R}^d$ is the usual *Euclidean distance* $\|\vec{u} - \vec{v}\| = \sqrt{(u_1 - v_1)^2 + \cdots + (u_d - v_d)^2}$. Finally, $\vec{u} \oplus \vec{v}$ is the vector from $\mathbb{R}^{d+d'}$ that concatenates $\vec{u} \in \mathbb{R}^d$ and $\vec{v} \in \mathbb{R}^{d'}$.

We focus on region-based embedding methods, whose regions are usually convex. A *region* X of $\mathbb{R}^d$ is a subset of $\mathbb{R}^d$. It is *convex* if for every $\vec{u}, \vec{v} \in X$ and $\lambda \in [0, 1]$, $(1 - \lambda)\vec{u} + \lambda\vec{v}$ is in X. Examples of convex regions are

- *convex cones*: for all $\vec{u}, \vec{v} \in X$, $\lambda, \mu \geq 0$, $\lambda\vec{u} + \mu\vec{v} \in X$;
- *boxes*: $X = \{\vec{x} \mid \vec{u} \leq \vec{x} \leq \vec{v}\}$ where $\vec{u}$ is the *lower corner* and $\vec{v}$ is the *upper corner* of the box;
- *balls*: an open (resp. closed) d-ball of radius ρ and center $\vec{x}$ is the set of all $\vec{y}$ such that $\|\vec{y} - \vec{x}\| < \rho$ (resp. $\leq \rho$).

Some embedding methods rely on *transformations* of $\mathbb{R}^d$, which are functions $f : \mathbb{R}^d \mapsto \mathbb{R}^d$. An *affine transformation* preserves convexity and parallelism and is defined by $f(\vec{x}) = A\vec{x} + \vec{b}$ with A an invertible matrix and $\vec{b} \in \mathbb{R}^d$. If A is the identity matrix, i.e. $f(\vec{x}) = \vec{x} + \vec{b}$, then f is a *translation*.

Embeddings We consider abstract notions of an embedding for a DL KB and an embedding method. We intentionally refrain from giving more precise definitions since the existing embeddings in the literature differ so much in the way in which they embed KBs. An *embedding* E is a function that maps the components of a KB (such as individual, concept and role names) into abstract structures associated with vector spaces (such as regions or vector transformations).

Definition 1 (Embedding method)**.** *An embedding method for $\mathcal{L}$ is an algorithm that given an ABox and a (possibly empty) $\mathcal{L}$-TBox, produces an embedding. We call embeddings generated by a given method M the M-embeddings.*

Embedding methods usually use *loss functions* that penalize, e.g., when regions associated to concepts or roles, or vectors associated with individuals, are not placed as expected. Embedding methods optimize the loss so that the embedding captures the KB knowledge. The loss function often uses a *margin parameter* which when less or equal to zero enforces that, for instance, the inclusion between regions is proper when the loss is zero. *Scoring functions* associate a score to facts or axioms, interpreted as how likely the fact or axiom is considered to be true. However, it is often difficult to have a fixed and pre-defined threshold for the scoring function which is rather only used to rank facts or axioms.

3 KB Embeddings and Their Semantics

Embeddings are usually used to assess facts or axioms (e.g., to predict plausible facts) but this can be done in different ways. Region-based embeddings come with a geometric-based semantics, but axioms' plausibility is also often evaluated using a scoring function, e.g., considering that an axiom is true if it gets a score above a threshold. This motivates the following definition of embedding semantics, which allows for considering various semantics for a given embedding.

Definition 2 (Embedding semantics)**.** *A semantics for an embedding method M is a function S_M, which given an M-embedding E and a language $\mathcal{L}$ returns a function $S_M(E, \mathcal{L})$ that maps each sentence in the language $\mathcal{L}$ to 1 (meaning true) or 0 (meaning false).*

Here, we focus on region-based embedding methods and their geometric-based semantics, hence we only consider one semantics for each method. Also, we consider one language per method. Thus, we may omit S and $\mathcal{L}$ and write $E \models_M \alpha$ for $S_M(E, \mathcal{L})(\alpha) = 1$ or $E \not\models_M \alpha$ for $S_M(E, \mathcal{L})(\alpha) = 0$.

We now briefly introduce the embedding methods we will consider in this paper, using the terminology and notation we introduced for embedding and semantics. Our focus is on KB embedding methods that can be applied to various DL languages but we also consider two KG embedding methods that are able to capture some patterns (cf. Table 3).

Convex geometric models (Gutiérrez-Basulto and Schockaert, 2018) This method applies to quasi-chained rules, which include in particular the description logic $\mathcal{ELHI}_\bot$ in normal form (concept inclusions of the form $A \sqsubseteq B$, $A_1 \sqcap A_2 \sqsubseteq B$, $\exists r^{(-)}.A \sqsubseteq B$ and $A \sqsubseteq \exists r^{(-)}.B$ with $A, A_i \in \mathsf{N_C} \cup \{\top\}$ and $B \in \mathsf{N_C} \cup \{\bot\}$ and role inclusions of the form $r \sqsubseteq s^{(-)}$, where $s^{(-)}$ can be a role name or its inverse). Each $a \in \mathsf{N_I}$ is embedded as a vector $E(a) \in \mathbb{R}^d$, each $A \in \mathsf{N_C}$ as a convex region $E(A) \subseteq \mathbb{R}^d$, and each $r \in \mathsf{N_R}$ as a convex region $E(r) \subseteq \mathbb{R}^{2d}$. The semantics of this method for $\mathcal{ELHI}_\bot$ is given by:

- $E \models_{conv} A(a)$ iff $E(a) \in E(A)$;
- $E \models_{conv} r(a, b)$ iff $E(a) \oplus E(b) \in E(r)$;
- $E \models_{conv} r \sqsubseteq s^{(-)}$ iff $E(r) \subseteq E(s^{(-)})$;
- $E \models_{conv} C \sqsubseteq D$ iff $E(C) \subseteq E(D)$;

where the embedding function E is extended to complex concept and role expressions in $\mathcal{ELHI}_\bot$ as follows (see (Bourgaux, Ozaki, and Pan, 2021) for a reference using a similar definition for a DL-Lite dialect):

- $E(\bot) := \emptyset$, $E(\top) := \mathbb{R}^d$;
- $E(r^-) := \{\vec{x} \oplus \vec{y} \mid \vec{x}, \vec{y} \in \mathbb{R}^d, \vec{y} \oplus \vec{x} \in E(r)\}$;
- $E(A_1 \sqcap A_2) := E(A_1) \cap E(A_2)$; and
- $E(\exists r^{(-)}.A) := \{\vec{x} \mid \vec{x} \in \mathbb{R}^d, \vec{x} \oplus \vec{y} \in E(r^{(-)}), \vec{y} \in E(A)\}$.

Al-cone models (Özçep, Leemhuis, and Wolter, 2020) Since in general the complement of a convex region may not be convex, when dealing with logics with negation, it is useful to consider convex regions which have a natural "complementary region" other than their actual complement. For this purpose, the authors of this method consider axis-aligned cones (al-cones), of the form $X_1 \times \cdots \times X_d$ with $X_i \in \{\mathbb{R}, \mathbb{R}_+, \mathbb{R}_-, \{0\}\}$. The method applies to (fragments of) $\mathcal{ALC}$ (which allows for concept inclusions using the $\sqcap$, $\sqcup$, $\neg$, $\exists$ and $\forall$ constructors). The authors consider *propositional* $\mathcal{ALC}$ which is a fragment of $\mathcal{ALC}$ that allows Boolean connectives ($\sqcap$, $\sqcup$, $\neg$) but disallows expressions containing roles. We call this fragment $\mathcal{ALC}_{\text{p}}$. The authors also consider the fragment of $\mathcal{ALC}$ that allows concept expressions with roles but limits the size of the expressions by a constant using a notion called *rank*. We denote this fragment with $\mathcal{ALC}_{\text{r}}$. Each $a \in \mathsf{N_I}$ is embedded as a vector $E(a) \in \mathbb{R}^d \setminus \{\vec{0}\}$, each $A \in \mathsf{N_C}$ as an al-cone $E(A)$, and each $r \in \mathsf{N_R}$ as a subset $E(r)$ of $\mathbb{R}^d \setminus \{\vec{0}\} \times \mathbb{R}^d \setminus \{\vec{0}\}$. The semantics of the al-cone embedding method for $\mathcal{ALC}$ is defined as:

- $E \models_{cone} A(a)$ iff $E(a) \in E(A)$;
- $E \models_{cone} r(a, b)$ iff $(E(a), E(b)) \in E(r)$;
- $E \models_{cone} C_1 \sqsubseteq C_2$ iff $E(C_1) \subseteq E(C_2)$;

where the embedding function E is extended to complex concepts C_1, C_2 as follows:

- $E(C_1 \sqcap C_2) := E(C_1) \cap E(C_2)$;
- $E(\neg C) := E(C)^o = \{\vec{x} \in \mathbb{R}^d \mid \forall \vec{y} \in E(C), \langle \vec{x}, \vec{y} \rangle \leq 0\}$ is the polar cone of $E(C)$;
- $E(C_1 \sqcup C_2) := E(\neg(\neg C_1 \sqcap \neg C_2))$;
- $E(\forall r.C)$ is the minimal al-cone containing $\{\vec{x} \mid (\vec{x}, \vec{y}) \in E(r) \Rightarrow \vec{y} \in E(C)\}$;
- $E(\exists r.C) := E(\neg \forall r.\neg C)$;
- $E(\top) := \mathbb{R}^d$; and $E(\bot) := \{\vec{0}\}$.

Remark 1. This semantics is such that it may be the case that $E \not\models_{cone} A(a)$ and $E \not\models_{cone} \neg A(a)$.

ELEm (Kulmanov et al., 2019) This method applies to a fragment of $\mathcal{EL}^{++}$ (Baader, Lutz, and Brandt, 2008) that corresponds to $\mathcal{ELO}_\bot$ (i.e. $\mathcal{EL}$ with nominals and $\bot$). Before being embedded, ABox assertions are transformed into TBox axioms using nominals as explained in Section 2.1 and the TBox is put in normal form. Each concept name or nominal C is embedded as an open d-ball $E(C) = \text{Ball}(C)$ represented by its center $c(C) \in \mathbb{R}^d$ and radius $\rho(C) \in \mathbb{R}$, and each $r \in \mathsf{N_R}$ as a vector $E(r) \in \mathbb{R}^d$. The top concept $\top$ is mapped to $\mathbb{R}^d$, that is, $\rho(\top) = \infty$. The semantics of the ELEm embedding method based on regions is defined for $\mathcal{ELO}_\bot$ axioms in normal form below.

- For (complex) concepts C and D different from $\bot$, $E \models_{elem} C \sqsubseteq D$ iff $\text{Ball}(C) \subseteq \text{Ball}(D)$ where
  - $\text{Ball}(C_1 \sqcap C_2) = \text{Ball}(C_1) \cap \text{Ball}(C_2)$,
  - $\text{Ball}(\exists r.C)$ is the ball with center $c(C) - E(r)$ and radius $\rho(C)$, i.e. $\text{Ball}(\exists r.C) = \{\vec{x} \mid \vec{x} + E(r) \in \text{Ball}(C)\}$.
- For concept inclusions with $\bot$ as right-hand side:
  - $E \models_{elem} A \sqsubseteq \bot$ iff $\rho(A) = 0$ (i.e. $\text{Ball}(A) = \emptyset$ since $\text{Ball}(A)$ is an open ball),
  - $E \models_{elem} \exists r.A \sqsubseteq \bot$ iff $\rho(A) = 0$, and
  - $E \models_{elem} A_1 \sqcap A_2 \sqsubseteq \bot$ iff $\text{Ball}(A_1) \cap \text{Ball}(A_2) \subseteq \emptyset$.

EmEL^{++} (Mondal, Bhatia, and Mutharaju, 2021) EmEL^{++} is similar to ELEm. The only difference is that Mondal, Bhatia, and Mutharaju (2021) additionally consider role inclusions and role composition (hence consider the fragment $\mathcal{ELHO}(\circ)_\bot$ of $\mathcal{EL}^{++}$), extending the semantics as follows:

- $E \models_{emel} r \sqsubseteq s$ iff $E(r) = E(s)$;
- $E \models_{emel} r_1 \circ r_2 \sqsubseteq s$ iff $E(r_1) + E(r_2) = E(s)$.

ELBE (Peng et al., 2022) ELBE is also similar to ELEm but uses boxes instead of balls, which has the advantage that the intersection of two boxes is still a box contrary to balls. Each concept name or nominal C is embedded as a box $E(C) = \text{Box}(C)$ represented by a pair of vectors $e_c(C)$ and $e_o(C)$ that represent the *center* and *offset* of the box. Specifically, the offset defines a non-negative real value for every dimension, such that $\vec{v} \in \text{Box}(C)$ iff $|\vec{v} - e_c(C)| \leq e_o(C)$. Each $r \in \mathsf{N_R}$ is embedded as a vector $E(r) \in \mathbb{R}^d$. We assume that the concept $\top$ is mapped to $\mathbb{R}^d$, that is, $e_o(\top) = \vec{\infty}$ (this is inspired by Kulmanov et al. (2019) but not explicit by Peng et al. (2022)). The semantics of the ELBE embedding method is defined as follows for $\mathcal{ELO}_\bot$ axioms in normal form.

- For (complex) concepts C and D different from $\bot$, $E \models_{elbe} C \sqsubseteq D$ iff $\text{Box}(C) \subseteq \text{Box}(D)$ where
  - $\text{Box}(C_1 \sqcap C_2) = \text{Box}(C_1) \cap \text{Box}(C_2)$,
  - $\text{Box}(\exists r.C) = \text{Box}(C) - E(r) = \{\vec{x} \mid \vec{x} + E(r) \in \text{Box}(C)\}$.
- For concept inclusions with $\bot$ as right-hand side:
  - $E \models_{elbe} A \sqsubseteq \bot$ iff $e_o(A) = \vec{0}$, and
  - $E \models_{elbe} \exists r.A \sqsubseteq \bot$ iff $e_o(A) = \vec{0}$.

BoxEL (Xiong et al., 2022) BoxEL also considers $\mathcal{ELO}_\bot$ and represents concepts as boxes, but represents roles through affine transformations instead of simple translations (in contrast with ELEm, EmEL^{++} and ELBE) in order to avoid that $A \sqsubseteq \exists r.B$ enforces that the volume of B is at least the one of A (i.e., to be able to represent many-to-one relations, that is, roles that are not inverse functional). Each $a \in \mathsf{N_I}$ is mapped to a vector $E(a) \in \mathbb{R}^d$. Each $A \in \mathsf{N_C}$ is embedded into a box $E(A) = \text{Box}(A)$, represented by two vectors from $\mathbb{R}^d$ which give its lower and upper corners. Each $r \in \mathsf{N_R}$ is embedded into an affine transformation $E(r) = T^r$ where $T^r(\vec{x}) = D^r\vec{x} + \vec{b^r}$ with D^r a diagonal matrix with non-negative entries and $\vec{b^r} \in \mathbb{R}^d$. The semantics of the BoxEL embedding method is given by geometric interpretations. Given an embedding E, the corresponding geometric interpretation is $\mathcal{I}_E = (\Delta^{\mathcal{I}_E}, \cdot^{\mathcal{I}_E})$ where $\Delta^{\mathcal{I}_E} = \mathbb{R}^d$ and

- for every $a \in \mathsf{N_I}$, $a^{\mathcal{I}_E} := E(a)$,
- for every $A \in \mathsf{N_C}$, $A^{\mathcal{I}_E} := \text{Box}(A)$, and
- for every $r \in \mathsf{N_R}$, $r^{\mathcal{I}_E} := \{(\vec{x}, \vec{y}) \mid T^r(\vec{x}) = \vec{y}\}$.

We then have that $E \models_{boxel} \alpha$ iff $\mathcal{I}_E \models \alpha$, where $\mathcal{I}_E$ is a standard DL interpretation.

BoxE (Abboud et al., 2020) A known issue with methods that embed roles using transformations of $\mathbb{R}^d$ is their inability to represent faithfully one-to-many relations, since transformations are functions. BoxE solved this issue by introducing so called 'bumps' to dynamically encode the relationship between entities and relations. It represents each relation r of arity n by a tuple of n boxes, $E(r) = (r^{(1)}, \ldots, r^{(n)})$, where each box $r^{(i)}$ is defined by two vectors that give its lower and upper corners, and each individual name $a \in \mathsf{N_I}$ by two vectors, $E(a) = (\vec{e_a}, \vec{b_a}) \in (\mathbb{R}^d)^2$, where $\vec{e_a}$ is the base position of the entity and $\vec{b_a}$ its translational bump. The semantics of the BoxE embedding method for the language $\mathcal{L}$ that consists of assertions as well as patterns from Table 3 except composition, is defined as follows:

- $E \models_{boxe} A(a)$ iff $\vec{e_a} \in A^{(1)}$;
- $E \models_{boxe} r(c, d)$ iff $\vec{e_c} + \vec{b_d} \in r^{(1)}$ and $\vec{e_d} + \vec{b_c} \in r^{(2)}$;
- $E \models_{boxe} r_1 \equiv r_2^-$ iff $r_1^{(1)} = r_2^{(2)}$ and $r_1^{(2)} = r_2^{(1)}$;
- $E \models_{boxe} r_1 \sqsubseteq r_2$ iff $r_1^{(1)} \subseteq r_2^{(1)}$ and $r_1^{(2)} \subseteq r_2^{(2)}$;
- $E \models_{boxe} r_1 \sqcap r_2 \sqsubseteq r_3$ iff $r_1^{(1)} \cap r_2^{(1)} \subseteq r_3^{(1)}$ and $r_1^{(2)} \cap r_2^{(2)} \subseteq r_3^{(2)}$;
- $E \models_{boxe} r_1 \sqsubseteq \neg r_2$ iff $r_1^{(1)} \cap r_2^{(1)} = \emptyset$ or $r_1^{(2)} \cap r_2^{(2)} = \emptyset$;
- $E \models_{boxe} r_1 \sqsubseteq \neg r_1^-$ iff $r_1^{(1)} \cap r_1^{(2)} = \emptyset$.

Box^{2}EL (Jackermeier, Chen, and Horrocks, 2024) This method applies to $\mathcal{ELHO}(\circ)_\perp$, the fragment of $\mathcal{EL}^{++}$ also considered by EmEL^{++}. It uses boxes and bumps, in line with Abboud et al. (2020). Similary to ELEm, ABox assertions are transformed into TBox axioms with nominals and the TBox is put in normal form. Each $A \in \mathsf{N_C}$ is represented by three vectors in $\mathbb{R}^d$, the first two being the lower and upper corners of a box $\mathrm{Box}(A)$ and the last one defining its bump: $E(A) = (\mathrm{Box}(A), \mathrm{Bump}(A))$. Each $a \in \mathsf{N_I}$ is represented by a vector $E(a) \in \mathbb{R}^d$ and nominal $\{a\}$ is mapped to $E(\{a\}) = (\mathrm{Box}(\{a\}), \mathrm{Bump}(\{a\}))$ where $\mathrm{Box}(\{a\})$ has volume 0 and is such that the lower and upper corners are equal to $E(a)$. Each $r \in \mathsf{N_R}$ is associated with two boxes $E(r) = (\mathrm{Head}(r), \mathrm{Tail}(r))$. The semantics of the Box^{2}EL embedding method is defined for $\mathcal{ELHO}(\circ)_\perp$ axioms in normal form as follows, where given a box B and a vector $\vec{v}$, $B + \vec{v} = \{\vec{x} + \vec{v} \mid \vec{x} \in B\}$ and similarly for $-$, and where a box with lower corner $\vec{l} = (l_1, \dots, l_d)$ and upper corner $\vec{u} = (u_1, \dots, u_d)$ is empty iff there exists i such that $l_i > u_i$.

- $E \models_{box2el} r_1 \sqsubseteq r_2$ iff $\mathrm{Head}(r_1) \subseteq \mathrm{Head}(r_2)$ and $\mathrm{Tail}(r_1) \subseteq \mathrm{Tail}(r_2)$;
- $E \models_{box2el} r_1 \circ r_2 \sqsubseteq s$ iff $\mathrm{Head}(r_1) \subseteq \mathrm{Head}(s)$ and $\mathrm{Tail}(r_2) \subseteq \mathrm{Tail}(s)$;
- $E \models_{box2el} A \sqsubseteq B$ iff $\mathrm{Box}(A) \subseteq \mathrm{Box}(B)$;
- $E \models_{box2el} A_1 \sqcap A_2 \sqsubseteq B$ iff $\mathrm{Box}(A_1) \cap \mathrm{Box}(A_2) \subseteq \mathrm{Box}(B)$;
- $E \models_{box2el} A \sqsubseteq \exists r.B$ iff $\mathrm{Box}(A) + \mathrm{Bump}(B) \subseteq \mathrm{Head}(r)$ and $\mathrm{Box}(B) + \mathrm{Bump}(A) \subseteq \mathrm{Tail}(r)$, and $\mathrm{Box}(A) \subseteq \emptyset$ if $\mathrm{Box}(B) = \emptyset$;
- $E \models_{box2el} \exists r.B \sqsubseteq A$ iff $\mathrm{Head}(r) - \mathrm{Bump}(B) \subseteq \mathrm{Box}(A)$;
- $E \models_{box2el} A \sqsubseteq \perp$ iff $\mathrm{Box}(A) = \emptyset$;
- $E \models_{box2el} A_1 \sqcap A_2 \sqsubseteq \perp$ iff $\mathrm{Box}(A_1) \cap \mathrm{Box}(A_2) = \emptyset$.

ExpressivE (Pavlovic and Sallinger, 2023) This KG embedding method embeds each $a \in \mathsf{N_I}$ as a vector $E(a) \in \mathbb{R}^d$ and each $r \in \mathsf{N_R}$ as an hyper-parallelogram in the virtual triple space $\mathbb{R}^{2d}$ (more precisely, r is mapped to three vectors from $\mathbb{R}^d$: a slope, a center and a width vector). The semantics of the ExpressivE embedding method for the language of role assertions and patterns from Table 3 is defined as follows.

- $E \models_{expr} r(a, b)$ iff $E(a) \oplus E(b) \in E(r)$;
- $E \models_{expr} r_1 \sqsubseteq r_1^-$ iff $E(r_1)$ is symmetric (i.e. is its own mirror image w.r.t. the identity line);
- $E \models_{expr} r_1 \equiv r_2^-$ iff $E(r_1)$ and $E(r_2)$ are mirror images of each other w.r.t. the identity line;
- $E \models_{expr} r_1 \sqsubseteq r_2$ iff $E(r_1) \subseteq E(r_2)$;
- $E \models_{expr} r_1 \sqcap r_2 \sqsubseteq r_3$ iff $E(r_1) \cap E(r_2) \subseteq E(r_3)$;
- $E \models_{expr} r_1 \sqsubseteq \neg r_2$ iff $E(r_1) \cap E(r_2) = \emptyset$;
- $E \models_{expr} r_1 \sqsubseteq \neg r_1^-$ iff $E(r_1)$ does not intersect with its mirror image;
- $E \models_{expr} r_1 \circ r_2 \sqsubseteq r_3$ iff $E(r_1 \circ r_2) \subseteq E(r_3)$ where $E(r_1 \circ r_2)$ is the compositionally defined convex region of r_1 and r_2, which is such that, for every $\vec{u}, \vec{v}, \vec{w} \in \mathbb{R}^d$, $\vec{u} \oplus \vec{v} \in E(r_1)$ and $\vec{v} \oplus \vec{w} \in E(r_2)$ iff $\vec{u} \oplus \vec{w} \in E(r_1 \circ r_2)$.

We point out that the works by Gutiérrez-Basulto and Schockaert (2018) and Özçep, Leemhuis, and Wolter (2020) have focused on theoretical aspects of their methods, without providing an implementation. The authors of the other embedding methods we describe above have provided implementations.

Other embedding methods have been designed for DLs in different contexts. For example, CosE (Li et al., 2022) embeds a DL-Lite$_{\mathsf{core}}$ TBox, seen as a KG with relations subClassOf and disjointWith to find plausible missing inclusion or disjointness between concepts. Closer to the methods we consider, TransOWL and TransROWL (d'Amato, Quatraro, and Fanizzi, 2021) were proposed for injecting background knowledge, which can be seen as a TBox. The basic idea is to modify the loss function by considering the facts (both positive and negative) inferred from those observed and the background knowledge. However, this method does not associate regions to concepts or roles, thus is out of the scope of this work. For a wider overview of possibly non-region based geometrical embeddings, see (Xiong et al., 2023).

4 Embedding Method Properties

We formulate theoretical properties for KB embeddings and embedding methods, show how they relate to each other and illustrate them on the embedding methods presented in Section 3. In this section, $\mathcal{L}$ denotes a DL language, M is an embedding method for $\mathcal{L}$, and S_M is a semantics for M.

Definition 3 (M-model)**.** *Let $\mathcal{A}$ be an ABox, $\mathcal{T}$ be a TBox in $\mathcal{L}$ and E be an M-embedding. The embedding E interpreted under S_M is an M-*model *of*

- *$\mathcal{A}$ if for every fact α of $\mathcal{A}$, $S_M(E, \mathcal{L})(\alpha) = 1$,*
- *$\mathcal{T}$ if for every axiom α of $\mathcal{T}$, $S_M(E, \mathcal{L})(\alpha) = 1$,*
- *$\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ if it is an M-model of $\mathcal{A}$ and $\mathcal{T}$.*

The existence of an M-*model* does not imply the existence of a model in the classical sense. That is, nothing in the definition of an M-model prevents $S_M(E, \mathcal{L})$ to assign to true inconsistent sets of axioms or to assign to false axioms that are entailed by axioms assigned to true.

Example 1. *Consider the (classically) unsatisfiable KB $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ with $\mathcal{T} = \{A \sqsubseteq \bot\}$ and $\mathcal{A} = \{A(a)\}$. Define an ELEm-embedding E of $\{A \sqsubseteq \bot, \{a\} \sqsubseteq A\}$ in $\mathbb{R}^2$ as follows: $E(\{a\}) = E(A) = \mathrm{Ball}(A)$ with center $c(A) = (0, 1)$ and radius $\rho(A) = 0$, i.e. $E(\{a\}) = E(A) = \emptyset$. It holds that $E \models_{elem} A \sqsubseteq \bot$ and $E \models_{elem} \{a\} \sqsubseteq A$ so E is an ELEm-model of $\mathcal{K}$. EmEL^{++} encounters the same problem since it translates assertions into TBox axioms then treats nominals as concepts so that they can be embedded to empty balls. BoxEL fixed this issue by mapping individuals to vectors.*

Conversely, non-existence of an M-model also does not imply non-existence of a model (in the classical sense).

Example 2. *As explained by Gutiérrez-Basulto and Schockaert (2018), the following KB does not have any convex geometric model while it is satisfiable: $\mathcal{T} = \{r_1 \sqsubseteq \neg r_2\}$ and $\mathcal{A} = \{r_1(a, b), r_1(b, a), r_2(a, a), r_2(b, b)\}$. Indeed, if E is a convex geometric model of $\mathcal{A}$, the following holds:*

- *$E(a) \oplus E(b) \in E(r_1)$ and $E(b) \oplus E(a) \in E(r_1)$ so that by convexity, $0.5(E(a) \oplus E(b)) + 0.5(E(b) \oplus E(a)) \in E(r_1)$;*
- *$E(a) \oplus E(a) \in E(r_2)$ and $E(b) \oplus E(b) \in E(r_2)$ so that by convexity, $0.5(E(a) \oplus E(a)) + 0.5(E(b) \oplus E(b)) \in E(r_2)$.*

Let $\vec{v} = 0.5E(a) + 0.5E(b)$. It holds that $\vec{v} \oplus \vec{v}$ is both in $E(r_1)$ and $E(r_2)$, so $E \not\models_{conv} r_1 \sqsubseteq \neg r_2$.

4.1 Soundness and Completeness

This section is concerned with the relationship between the existence of an M-model and that of a classical model.

Property 1 (Embedding method soundness)**.** *We say that M under S_M is* sound for $\mathcal{L}$ *if the existence of an M-model (under S_M) for a KB $\mathcal{K}$ in $\mathcal{L}$ implies that $\mathcal{K}$ is satisfiable.*

Property 2 (Embedding method completeness)**.** *We say that M under S_M is* complete for $\mathcal{L}$ *if for every satisfiable KB $\mathcal{K}$ in $\mathcal{L}$, there is an M-model (under S_M) for $\mathcal{K}$.*

Example 3. *Corollary 1 in (Gutiérrez-Basulto and Schockaert, 2018) states that embedding methods that produce convex geometric models are sound and complete for the language of quasi-chained rules (hence in particular for $\mathcal{ELHI}_\bot$ in normal form), and Proposition 2 in (Özçep, Leemhuis, and Wolter, 2020) states that methods that produce al-cones models are sound and complete for $\mathcal{ALC}_p$.*

As recalled in Example 2, embedding methods that produce convex geometric models are not complete for languages with role disjointness, under a semantics where role disjointness means disjointness of the role embeddings. Example 4 shows that BoxE (which does not fall into this class) is also incomplete for languages with role disjointness.

Example 4. *Consider the satisfiable KB $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ with $\mathcal{A} = \{r(a, b), s(a, c), r(d, c), s(d, b)\}$ and $\mathcal{T} = \{r \sqsubseteq \neg s\}$. Assume for a contradiction that there exists a BoxE-model E of $\mathcal{K}$. Recall that E maps each role r to two boxes, one for the "head", denoted $r^{(1)}$, and one for the "tail", denoted $r^{(2)}$. Also, recall that each box $r^{(i)}$ is represented by its lower and upper corners, denoted $\vec{l_{r^{(i)}}}$ and $\vec{u_{r^{(i)}}}$ respectively. Moreover, a point $\vec{e}$ is in a box $r^{(i)}$ if it is between its lower and upper corners, in symbols, $\vec{l_{r^{(i)}}} \leq \vec{e} \leq \vec{u_{r^{(i)}}}$.*

Since $E \models_{boxe} r \sqsubseteq \neg s$, then $r^{(1)} \cap s^{(1)} = \emptyset$ or $r^{(2)} \cap s^{(2)} = \emptyset$. Assume $r^{(1)} \cap s^{(1)} = \emptyset$ (the argument for the case where $r^{(2)} \cap s^{(2)} = \emptyset$ is analogous). Given a vector $\vec{v}$, denote by $\vec{v}[k]$ the value of $\vec{v}$ at position k. As $r^{(1)} \cap s^{(1)} = \emptyset$, there is a dimension j such that

$$\vec{u_{r^{(1)}}}[j] < \vec{l_{s^{(1)}}}[j] \text{ or } \vec{u_{s^{(1)}}}[j] < \vec{l_{r^{(1)}}}[j].$$

Suppose $\vec{u_{r^{(1)}}}[j] < \vec{l_{s^{(1)}}}[j]$. As $E \models_{boxe} r(a, b)$ and $E \models_{boxe} s(a, c)$, it must be the case that

$$\vec{e_a}[j] + \vec{b_b}[j] \leq \vec{u_{r^{(1)}}}[j] < \vec{l_{s^{(1)}}}[j] \leq \vec{e_a}[j] + \vec{b_c}[j]$$

which implies that $\vec{b_b}[j] < \vec{b_c}[j]$. Now, as $E \models_{boxe} r(d, c)$ and $E \models_{boxe} s(d, b)$, we obtain

$$\vec{e_d}[j] + \vec{b_c}[j] \leq \vec{u_{r^{(1)}}}[j] < \vec{l_{s^{(1)}}}[j] \leq \vec{e_d}[j] + \vec{b_b}[j]$$

which implies that $\vec{b_c}[j] < \vec{b_b}[j]$, contradicting $\vec{b_b}[j] < \vec{b_c}[j]$. The case $\vec{u_{s^{(1)}}}[j] < \vec{l_{r^{(1)}}}[j]$ can be proved analogously.

In the literature, it is common to consider an alternative meaning for soundness, which intuitively links the existence of an embedding with loss 0 and KB satisfiability. A loss function associated with an embedding method M can be seen as a function loss that takes as input a KB $\mathcal{K}$ and an M-embedding E of $\mathcal{K}$ and returns a number.

Property 3 (Embedding method soundness based on loss)**.** *If M has a loss function* loss*, we say that M is* sound for $\mathcal{L}$ w.r.t. the loss function *if the existence of an M-embedding E of a KB $\mathcal{K}$ in $\mathcal{L}$ such that $\mathrm{loss}(\mathcal{K}, E) = 0$ implies that $\mathcal{K}$ is satisfiable.*

Example 5. *Example 1 shows that ELEm is not sound. Moreover, it also shows that ELEm is not sound w.r.t. the loss function defined in (Kulmanov et al., 2019). If the margin parameter γ is equal to 0 since $E(\{a\})$ and $E(A)$ have the same center that lies on the unity sphere and the same radius 0, the loss of the axiom $\{a\} \sqsubseteq A$ given by $\max(0, \|c(\{a\}) - c(A)\| + \rho(\{a\}) - \rho(A) - \gamma) + |\|c(\{a\})\| - 1| + |\|c(A)\| - 1|$ is equal to 0, and the loss of the axiom $A \sqsubseteq \bot$ given by $\rho(A)$ is equal to 0. Hence, $\mathrm{loss}(\mathcal{K}, E) = 0$. For Theorem 1 in (Kulmanov et al., 2019) to hold, γ should be strictly negative, rather than ≤ 0.[1] This however prevents ELEm to embed equivalent concepts, such as $\{A \sqsubseteq B, B \sqsubseteq A\}$, with loss 0.*

Example 6 illustrates the difference between soundness w.r.t. the loss function and soundness as in Property 1.

Example 6. *Theorem 1 in (Jackermeier, Chen, and Horrocks, 2024) shows that Box^{2}EL is sound w.r.t. the loss*

[1] With a minor fix in (Kulmanov et al., 2019, Equation 2): one of the subterms of the loss term for $A_1 \sqcap A_2 \sqsubseteq B$ is $\max(0, \|c(A_1) - c(B)\| - \rho(A_1) - \gamma)$. The following subterm is written as $\max(0, \|c(A_2) - c(B)\| - \rho(A_1) - \gamma)$ while it should be analogous to the one before, using $\rho(A_2)$ instead of $\rho(A_1)$.

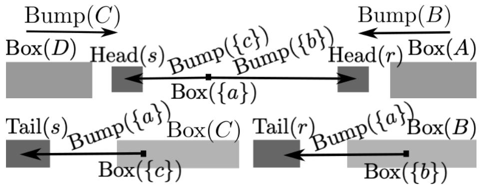

$$\exists r.B \sqsubseteq A: \quad \mathrm{Head}(r) - \mathrm{Bump}(B) \subseteq \mathrm{Box}(A)$$
$$\exists s.C \sqsubseteq D: \quad \mathrm{Head}(s) - \mathrm{Bump}(C) \subseteq \mathrm{Box}(D)$$
$$A \sqcap D \sqsubseteq \bot: \quad \mathrm{Box}(A) \cap \mathrm{Box}(D) = \emptyset$$
$$r(a,b): \quad \begin{Bmatrix} \mathrm{Box}(\{a\}) + \mathrm{Bump}(\{b\}) \subseteq \mathrm{Head}(r) \\ \mathrm{Box}(\{b\}) + \mathrm{Bump}(\{a\}) \subseteq \mathrm{Tail}(r) \end{Bmatrix}$$
$$s(a,c): \quad \begin{Bmatrix} \mathrm{Box}(\{a\}) + \mathrm{Bump}(\{c\}) \subseteq \mathrm{Head}(s) \\ \mathrm{Box}(\{c\}) + \mathrm{Bump}(\{a\}) \subseteq \mathrm{Tail}(s) \end{Bmatrix}$$
$$B(b): \mathrm{Box}(\{b\}) \subseteq \mathrm{Box}(B) \quad C(c): \mathrm{Box}(\{c\}) \subseteq \mathrm{Box}(C)$$

Figure 2: Box^{2}EL-embedding.

function. However, Box^{2}EL is not sound. Indeed, the loss function of Box^{2}EL is such that in models of loss 0, all bumps are equal to $\vec{0}$, while for soundness, we consider also models with non-zero bumps. To illustrate this, consider $\mathcal{T} = \{\exists r.B \sqsubseteq A, \exists s.C \sqsubseteq D, A \sqcap D \sqsubseteq \bot\}$ and $\mathcal{A} = \{r(a,b), s(a,c), B(b), C(c)\}$. The KB $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is unsatisfiable. However, the Box^{2}EL-embedding depicted in Figure 2 is a Box^{2}EL-model of $\mathcal{K}$. Note that Box^{2}EL-embeddings with loss 0 have an undesirable behaviour: since all bumps are $\vec{0}$, $E \models_{box2el} A \sqsubseteq \exists r.B$ and $E \models_{box2el} C \sqsubseteq \exists r.D$ imply $E \models_{box2el} A \sqsubseteq \exists r.D$ and $E \models_{box2el} C \sqsubseteq \exists r.B$.

4.2 Entailment Closure and Faithfulness

Since M-models come with very few guarantees on what they assign to true besides the KB itself, additional properties can be required on the M-models. Entailment closure guarantees that all consequences of the KB are assigned to true.

Definition 4 (Entailment closure in an M-model)**.** *Let $\mathcal{T}$ be a TBox in $\mathcal{L}$ and $\mathcal{A}$ be an ABox such that $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is satisfiable. We say that an M-model E of $\mathcal{K}$ is*

- TBox-entailed *for $\mathcal{L}$ if for every TBox axiom α in $\mathcal{L}$ that is entailed by $\mathcal{K}$, $E \models_M \alpha$;*
- ABox-entailed *if for every assertion α that is entailed by $\mathcal{K}$, $E \models_M \alpha$;*
- KB-entailed *if it is TBox-entailed and ABox-entailed.*

A slight modification of Example 6 provides a Box^{2}EL-model that is not ABox-entailed.

Example 7. *Consider $\mathcal{T} = \{\exists r.B \sqsubseteq A, \exists s.C \sqsubseteq D\}$ and $\mathcal{A} = \{r(a,b), s(a,c), B(b), C(c)\}$. The KB $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is satisfiable and $\mathcal{K}$ entails $A(a)$ and $D(a)$. However, the Box^{2}EL-model of $\mathcal{K}$ in Figure 2 does not satisfy $A(a)$, $D(a)$.*

Entailment closure does not prevent the embedding semantics to assign to true axioms that are, e.g., not consistent with the KB. The notions of weak and strong faithfulness have been proposed in the literature and address this issue.

Definition 5 (Weak faithfulness of an M-model (adapted from (Özçep, Leemhuis, and Wolter, 2020)))**.** *Let $\mathcal{T}$ be a TBox in $\mathcal{L}$ and let $\mathcal{A}$ be an ABox such that $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is satisfiable. We say that an M-model E of $\mathcal{K}$ is*

- weakly TBox-faithful *for $\mathcal{L}$ if for every TBox axiom α in $\mathcal{L}$, $E \models_M \alpha$ implies that α is consistent with $\mathcal{K}$;*
- weakly ABox-faithful *if for every assertion α, $E \models_M \alpha$ implies that α is consistent with $\mathcal{K}$;*
- weakly KB-faithful *if it is weakly TBox-faithful and weakly ABox-faithful.*

Example 8 shows that some KBs may have only ELEm- or EmEL^{++}-models that are not weakly ABox-faithful.

Example 8. *Let $\mathcal{T} = \{B \sqcap C \sqsubseteq \bot\}$ and $\mathcal{A} = \{r(a,b), r(a,c), B(b), C(c)\}$. For every M-model E of $\mathcal{A}$ and $\mathcal{T}$ with $M \in \{$ELEm, EmEL$^{++}\}$, $E \models_M C(a)$ and $E \models_M B(a)$. Indeed, since $E \models_M \{a\} \sqsubseteq \exists r.\{b\}$, it holds that $E(\{a\}) \subseteq E(\{b\}) - E(r)$ and similarly, $E(\{a\}) \subseteq E(\{c\}) - E(r)$. Since $E(\{b\}) \subseteq E(B)$, and $E(\{c\}) \subseteq E(C)$, it follows that $E(\{a\}) + E(r) \subseteq E(B) \cap E(C) = \emptyset$, i.e., $E(\{a\}) = \emptyset$ is included in every region of $\mathbb{R}^d$.*

A stronger condition than weak faithfulness ensures that models satisfy *only* the KB consequences.

Definition 6 (Strong faithfulness of an M-model (adapted from (Özçep, Leemhuis, and Wolter, 2020)))**.** *Let $\mathcal{T}$ be a TBox in $\mathcal{L}$ and $\mathcal{A}$ an ABox such that $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$ is satisfiable. We say that an M-model E of $\mathcal{K}$ is*

- strongly TBox-faithful *for $\mathcal{L}$ if for every TBox axiom α in $\mathcal{L}$, $E \models_M \alpha$ implies that α is entailed by $\mathcal{K}$;*
- strongly ABox-faithful *if, for every assertion α, $E \models_M \alpha$ implies that α is entailed by $\mathcal{K}$;*
- strongly KB-faithful *if it is strongly TBox-faithful and strongly ABox-faithful.*

Example 9 illustrates that some KBs may have only Box^{2}EL-models that are not strongly TBox-faithful.

Example 9. *Consider $\mathcal{T} = \{r_1 \circ r_2 \sqsubseteq r_3, \exists r_3.C \sqsubseteq D\}$. Let E be a Box^{2}EL-model of $\mathcal{T}$. Since $E \models_{box2el} r_1 \circ r_2 \sqsubseteq r_3$, then $\mathrm{Head}(r_1) \subseteq \mathrm{Head}(r_3)$, and since $E \models_{box2el} \exists r_3.C \sqsubseteq D$, then $\mathrm{Head}(r_3) - \mathrm{Bump}(C) \subseteq \mathrm{Box}(D)$. It follows that $\mathrm{Head}(r_1) - \mathrm{Bump}(C) \subseteq \mathrm{Box}(D)$, so $E \models_{box2el} \exists r_1.C \sqsubseteq D$. However, $\mathcal{T} \not\models \exists r_1.C \sqsubseteq D$.*

Example 10 shows that some KBs may have only ELEm- or EmEL^{++}-models that are not strongly TBox-faithful.

Example 10. *Let $\mathcal{T} = \{\exists r.C \sqsubseteq A, \exists r.D \sqsubseteq B, A \sqcap B \sqsubseteq \bot\}$. For every M-model E of $\mathcal{T}$ with $M \in \{$ELEm, EmEL$^{++}\}$, $E(C) - E(r) \subseteq E(A)$, $E(D) - E(r) \subseteq E(B)$ and $E(A) \cap E(B) \subseteq \emptyset$. Hence $E(C) \cap E(D) \subseteq \emptyset$. It follows that $E \models_M C \sqcap D \sqsubseteq \bot$ while $\mathcal{T} \not\models C \sqcap D \sqsubseteq \bot$.*

Remark 2. The dimension of the embedding space $\mathbb{R}^d$ may have a strong impact on strong TBox faithfulness. Indeed, if $\mathcal{L}$ is a language allowing for concept intersections, and M is an embedding method which maps concepts to convex regions, $\bot$ to $\emptyset$, and interprets the conjunction of concepts as the intersection of their embeddings, then for every $k > d+1$, the TBox $\mathcal{T} = \{C_1 \sqcap ... \sqcap C_k \sqsubseteq \bot\}$ is such that no M-model of $\mathcal{T}$ is strongly TBox-faithful. Indeed, if E is an M-model of $\mathcal{T}$, then $\bigcap_{i=1}^{k} E(C_i) = \emptyset$ and Helly's theorem (Helly, 1923) states that if $X_1, \ldots, X_k$ are convex regions in $\mathbb{R}^d$, with $k > d$, and each $d+1$ among these regions have a non-empty intersection, it holds that $\bigcap_{i=1}^{k} X_i \neq \emptyset$. Hence, there must be some $\{C_{i_1}, \ldots, C_{i_{d+1}}\} \subsetneq \{C_1, \ldots, C_k\}$ such that $E \models_M C_{i_1} \sqcap \ldots \sqcap C_{i_{d+1}} \sqsubseteq \bot$.

Entailment closure and strong faithfulness together guarantee that an M-model behaves as a canonical model.

Proposition 1. *Let $\mathcal{T}$ be an $\mathcal{L}$-TBox, $\mathcal{A}$ an ABox and E an M-model of $\mathcal{K} = \mathcal{T} \cup \mathcal{A}$. Then the following holds.*

- *If E is TBox-entailed and strongly TBox-faithful then for every TBox axiom α in $\mathcal{L}$, $\mathcal{K} \models \alpha$ iff $E \models_M \alpha$.*
- *If E is ABox-entailed and strongly ABox-faithful then for every assertion α, $\mathcal{K} \models \alpha$ iff $E \models_M \alpha$.*

To study embedding methods w.r.t. entailment closure and faithfulness, we define the following properties.

Property 4 (Ability)**.** *Let $\mathcal{Y} \in \{$TBox, ABox, KB$\}$. We say that M under S_M is able to be (weakly/strongly) $\mathcal{Y}$-faithful for a language $\mathcal{L}$ if for every satisfiable $\mathcal{L}$-KB $\mathcal{K}$, there exists an M-model E of $\mathcal{K}$ such that E interpreted under S_M is (weakly/strongly) $\mathcal{Y}$-faithful. The $\mathcal{Y}$-entailed ability is defined as expected.*

Property 5 (Guarantee)**.** *Let $\mathcal{Y} \in \{$TBox, ABox, KB$\}$. We say that M under S_M is guaranteed to be (weakly/strongly) $\mathcal{Y}$-faithful for a language $\mathcal{L}$ if, for every satisfiable KB $\mathcal{K}$, M always produces an M-model E of $\mathcal{K}$ such that E interpreted under S_M is (weakly/strongly) $\mathcal{Y}$-faithful. The $\mathcal{Y}$-entailed guarantee is defined as expected.*

Proposition 2. *If M is guaranteed to be both strongly (resp. weakly) TBox-faithful and strongly (resp. weakly) ABox-faithful for $\mathcal{L}$ then it is guaranteed to be strongly (resp. weakly) KB-faithful for $\mathcal{L}$. The same holds when considering the entailment closure guarantee property.*

This does not hold for the *ability* properties: e.g., the existence of a TBox-entailed M-model and the existence of an ABox-entailed M-model do not imply the existence of a KB-entailed M-model.

4.3 Expressiveness

We extend the notion of full expressiveness (Kazemi and Poole, 2018), a well-known characteristic considered for KG embeddings, to languages that include TBox axioms.

Property 6 (Full Expressiveness)**.** *M under S_M is*

- fully TBox-expressive *for $\mathcal{L}$ if for every two $\mathcal{L}$-TBoxes $\mathcal{T}, \mathcal{T}'$, with $\mathcal{T}$ satisfiable and $\mathcal{T}'$ disjoint from the deductive closure of $\mathcal{T}$, there exists an M-model E of $\mathcal{T}$ such that $S_M(E, \mathcal{L})(\alpha) = 0$ for all $\alpha \in \mathcal{T}'$;*
- fully ABox-expressive *if for every two ABoxes $\mathcal{A}, \mathcal{A}'$, with $\mathcal{A}'$ being disjoint from $\mathcal{A}$, there exists an M-model E of $\mathcal{A}$ such that $S_M(E, \mathcal{L})(\alpha) = 0$ for all $\alpha \in \mathcal{A}'$.*

Full TBox-expressiveness is extended for KBs as expected.

Full ABox-expressiveness coincides with the notion of full expressiveness from the KG embedding literature, and is tightly related to strong ABox-faithfulness.

Proposition 3. *M under S_M is fully ABox-expressive iff for any ABox $\mathcal{A}$ there is an M-model E of $\mathcal{A}$ interpreted under S_M that is strongly ABox-faithful.*

In the KG literature, authors often consider the ability of *capturing patterns* from Table 3 (Abboud et al., 2020; Pavlovic and Sallinger, 2023). They distinguish the ability to capture a single pattern or to capture jointly several patterns (possibly of different kinds). Indeed, some methods are able to produce an embedding that captures a pattern but not multiple ones, even of the same type (Abboud et al., 2020, Table 1). It follows from the form of the patterns that all sets of patterns are satisfiable (if no facts need to be considered).

Definition 7 (Capturing patterns (adapted from (Abboud et al., 2020)))**.** *Let $\mathcal{L}$ be a language of patterns and let E be an M-embedding. E interpreted under S_M*

- *captures* exactly a pattern *$\phi \in \mathcal{L}$ if $S_M(E, \mathcal{L})(\phi) = 1$;*
- *captures exactly a* set of patterns *$\mathcal{S} = \{\phi_1, \ldots, \phi_n\} \subseteq \mathcal{L}$ if it captures exactly ϕ_i, for all $1 \leq i \leq n$;*
- *captures* exclusively *a set of patterns $\mathcal{S}$ if for every pattern ϕ in $\mathcal{L}$, $S_M(E, \mathcal{L})(\phi) = 1$ only if $\mathcal{S} \models \phi$.*

Property 7 (Ability to capture (adapted from (Abboud et al., 2020)))**.** *We say that M under S_M is able to capture (exactly/exclusively) $\mathcal{L}$ if for any finite set of patterns $\mathcal{S}$ expressed in $\mathcal{L}$, there exists an M-embedding interpreted under S_M that captures (exactly/exclusively) $\mathcal{S}$.*

We relate this property with strong faithfulness ability.

Proposition 4. *If $\mathcal{L}$ is a language of patterns, then M is able to capture exactly and exclusively $\mathcal{L}$ iff for any finite set of patterns $\mathcal{S}$ expressed in $\mathcal{L}$, there exists a strongly TBox-faithful M-model of $\mathcal{S}$.*

Example 11. *Let $\mathcal{L}$ be the language of exclusion patterns built on $\mathsf{N_R} = \{r_1, r_2\}$ (i.e. $\mathcal{L} = \{r_1 \sqsubseteq \neg r_2, r_2 \sqsubseteq \neg r_1\}$). Theorem 5.1 in Pavlovic and Sallinger (2023) shows that ExpressivE is fully ABox-expressive. It can be easily checked that ExpressivE is fully TBox-expressive for $\mathcal{L}$ (see (Pavlovic and Sallinger, 2023, Theorem 5.2)). However, ExpressivE is* not *fully KB-expressive for this language because it is not complete for $\mathcal{L}$ (see Example 2, which can be instantiated for ExpressivE since the hyper-parallelograms are convex).*[2] *In the same way, Theorem 5.1 in (Abboud et al., 2020) shows that BoxE is fully ABox-expressive and Theorem 5.3 shows that it is fully TBox-expressive for $\mathcal{L}$ while, by Example 4, BoxE is not complete thus* not *fully KB-expressive for $\mathcal{L}$.*

Some authors also consider *knowledge injection* (Benedikt et al., 2020), which broadly refers to the task of incorporating explicit (pre-defined) knowledge expressed as rules or constraints into a machine learning model. This can be achieved by constraining the training, the output, or the model itself with such patterns required to hold in the model. Abboud et al. (2020) establish that their embeddings can be modified so as to provably ensure that they satisfy some patterns (among a restricted class of patterns).

Embedding method TBox language	Convex $\mathcal{ELHI}_\bot^1$	Al-cone $\mathcal{ALC}_\mathrm{r}$	ELEm $\mathcal{ELO}_\bot^1$	EmEL^{++} $\mathcal{ELHO}(\circ)_\bot^1$	ELBE $\mathcal{ELO}_\bot^1$	BoxEL $\mathcal{ELO}_\bot^1$	Box^{2}EL $\mathcal{ELHO}(\circ)_\bot^1$	BoxE patterns^2	ExpressivE patterns
Soundness	✓	✓	✗	✗	✗	✓	✗	✓	✓
Completeness	✓	✓	✗	✗	✗	✗	✗	✗	✗
∀ ABox-Entailed	✓	✓‡	✗	✗	✗	✗	✗	✗	✗
∀ TBox-Entailed	✓	✓‡	✗	✗	✗	✗	✗	✗	✗
∃ Weak ABox-Faithful.	✓	✓	✗	✗	✗	✗	✗	✗	✗
∃ Weak TBox-Faithful.	✓	✓‡	✗	✗	✗	✗	✗	✗	✗
∃ Weak KB-Faithful.	✓	✓‡	✗	✗	✗	✗	✗	✗	✗
∀ Weak ABox-Faithful.	✓	✓	✗	✗	✗	✗	✗	✗	✗
∀ Weak TBox-Faithful.	✓	?	✗	✗	✗	✗	✗	✗	✗
∃ Strong ABox-Faithful.	✓†	✓	✗	✗	✗	✗	✗	✗	✗
∃ Strong TBox-Faithful.	✓†	✓‡	✗	✗	✗	✗	✗	✗	✗
∃ Strong KB-Faithful.	✓†	✓‡	✗	✗	✗	✗	✗	✗	✗
∀ Strong ABox-Faithful.	✗	✗	✗	✗	✗	✗	✗	✗	✗
∀ Strong TBox-Faithful.	✗	✗	✗	✗	✗	✗	✗	✗	✗
Full ABox Expressiveness	✓	✓	✗	✗	✗	✗	✓	✓	✓
Full TBox Expressiveness	✓†	✓‡	✗	✗	✗	✗	✗	✓	✓°

Table 4: Properties of KB embedding methods. [1] in normal form. [2] without composition. For † cases, results shown for $\mathcal{ELH}$ but we conjecture they also hold for $\mathcal{ELHI}_\bot$. For ‡ cases, results are for $\mathcal{ALC}_\mathrm{p}$. The ° result is for the language of *positive* patterns (without negation). Since all the languages considered are finite, ∃ ABox-, TBox-, KB-entailed coincide with completeness, and full KB-expressiveness coincides with ∃ strong KB-faithfulness (cf. Figure 3). By Proposition 2, for X ∈ {Entailed, Strong Faith., Weak Faith.}, ∀ KB-X holds iff ∀ ABox-X and ∀ TBox-X hold.

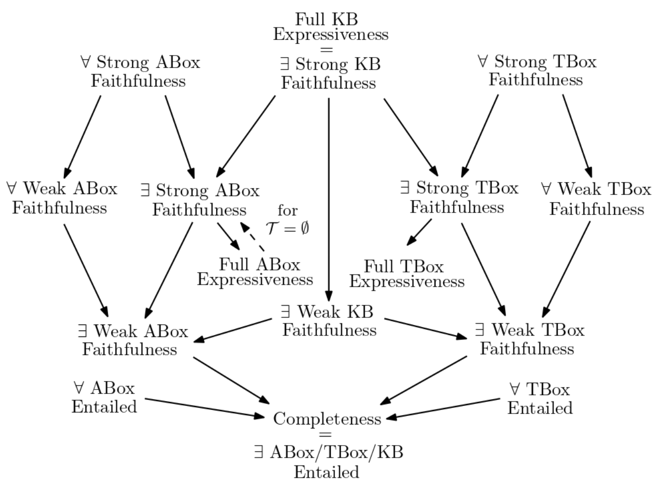

Figure 3: Relationships between the properties when the DL language is *finite*. An arrow from X to Y indicates that X implies Y.

4.4 Relationships Between Properties

We now briefly discuss the relationships between the properties, considering two cases: one for the general case (with possibly *infinite* languages) (Figure 1), and one for the special case of *finite* languages (Figure 3). For readability, we omit ∀ KB Entailed, ∀ Strong KB Faithfulness and ∀ Weak KB Faithfulness since they are equivalent to the conjunction of the ABox and TBox versions of the properties (Proposition 2). Note that many properties imply completeness because their definitions assume the existence of an M-model, and that guarantees imply abilities also because the guarantee definition requires the existence of an M-model. The other relationships are more informative (e.g., strong faithfulness implies weak faithfulness, and strong KB-faithfulness implies full KB-expressiveness).

Theorem 1. *The relationships between the properties of embedding methods shown in Figure 1 hold.*

When the language is finite, we observe that some properties coincide. First, being able to be ABox-, TBox- and KB-entailed is equivalent to being complete. Indeed, in this case, the deductive closure of the KB is finite so by completeness, one can find an M-model that satisfies every consequence of the KB. Second, being able to be strongly KB-faithful coincides with being fully KB-expressive for a similar reason: one can use the deductive closure to obtain a strongly KB-faithful M-model when M is fully KB-expressive.

Theorem 2. *For finite languages, the relationships between the properties of embedding methods in Figure 3 hold.*

5 Properties of Selected Methods

Table 4 shows which of the KB embedding methods of Section 3 satisfy the properties introduced in Section 4. Since we consider KBs in normal form and finite sets $\mathsf{N_I}$, $\mathsf{N_C}$, and $\mathsf{N_R}$, all languages are finite so we only consider properties of Figure 3. This comparison is not intended to be used to claim that some embedding methods are better than others based on the number of properties they satisfy. Our goal here is only to better understand the theoretical properties of these methods. Indeed, recall that not all methods apply to the same languages, so they cannot be directly compared.

[2] Though, the authors also consider a weaker semantics where pattern satisfaction is defined w.r.t. grounded pattern instances.

Also, as mentioned earlier, those that satisfy more properties, namely, Convex and Al-cone, are not implemented.

Moreover, depending on the use case, some properties may not be desirable. In particular, strong ABox-faithful M-models are unable to predict new plausible facts that are not entailed by the KB. Actually, we should often aim for embedding methods that are sound, complete and guaranteed to be strongly TBox-faithful and weakly ABox-faithful. Indeed, M-models that are strongly TBox-faithful and weakly ABox-faithful give formal guarantees that the TBox part from the source KB is respected, while still allowing downstream tasks such as link prediction to be performed based on the data coming from the ABox.

Most methods we consider either do not cover or fail to represent role composition, mutual exclusion, and axioms of the form $\exists r.C \sqsubseteq \bot$. E.g., BoxE and ExpressivE are not fully KB expressive for languages with mutual exclusion (Example 11). Box^{2}EL and EmEL^{++} fail to represent role composition (see Example 9 and the definition of EmEL^{++}, which implies that if $E \models_{emel} r_1 \circ r_2 \sqsubseteq s$ then $E \models_{emel} r_2 \circ r_1 \sqsubseteq s$). Also, ELEm, EmEL^{++}, and ELBE cannot precisely handle the concept inclusion $\exists r.C \sqsubseteq \bot$. These methods approximate it by $C \sqsubseteq \bot$, which implies $\exists r.C \sqsubseteq \bot$ but is not equivalent.

6 Conclusion and Perspectives

In this work, we examine recent region-based KB embedding methods through the lens of the properties of their geometric-based semantics. Our framework provides a common vocabulary and clarifies relationships between properties of KB embeddings (many of them already considered in the literature). It can be used to guide analysis of new methods and facilitate comparisons between existing and future embedding methods. In particular, while several theoretical properties have been considered relevant, such as those related to faithfulness, in practice, there are no implementations that satisfy them. Hence a novel practical embedding method that would satisfy, e.g., soundness and completeness would already offer more theoretical guarantees than existing ones. The main difficulties encountered for fulfilling the properties are related to the ability of representing role disjointness (together with facts!) and the bottom concept. This calls for some research effort since many natural constraints involve these constructs. For example, Wikidata has hundreds of "conflict with" constraints, which correspond to disjointness axioms between (complex) DL concepts.

Recent works in the KG literature focus on query answering, where the task is not only to rank facts but expressions in a richer query language (Hamilton et al., 2018; Ren and Leskovec, 2020; Ren, Hu, and Leskovec, 2020; Zhang et al., 2021; Bai et al., 2023). Note that these works consider KGs rather than KBs. An exception is the work by Imenes, Guimarães, and Ozaki (2023) which targets DL-Lite KBs by performing query rewriting and then querying the ABox embedding. We could extend several of our properties to consider queries, for example define weak or strong query-faithfulness for some query language. This would require to extend the embedding method semantics to evaluate queries.

Acknowledgements

This work was supported by the "PHC AURORA" programme (projects 49698VH and 341302), funded by the French Ministry for Europe and Foreign Affairs, the French Ministry for Higher Education and Research, and the Research Council of Norway, as well as the OLEARN project (316022) funded by the Research Council of Norway.

References

Abboud, R.; Ceylan, İ. İ.; Lukasiewicz, T.; and Salvatori, T. 2020. BoxE: A box embedding model for knowledge base completion. In *Proceedings of NeurIPS*.

Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press, 1st edition.

Baader, F.; Lutz, C.; and Brandt, S. 2008. Pushing the EL envelope further. In *Proceedings of the Fourth OWLED Workshop on OWL: Experiences and Directions*.

Bai, Y.; Lv, X.; Li, J.; and Hou, L. 2023. Answering complex logical queries on knowledge graphs via query computation tree optimization. In *Proceedings of ICML*.

Benedikt, M.; Kersting, K.; Kolaitis, P. G.; and Neider, D. 2020. Logic and Learning (Dagstuhl Seminar 19361). *Dagstuhl Reports* 9(9):1–22.

Bourgaux, C.; Guimarães, R.; Koudijs, R.; Lacerda, V.; and Ozaki, A. 2024. Knowledge base embeddings: Semantics and theoretical properties. arXiv:2408.04913 [cs.AI].

Bourgaux, C.; Ozaki, A.; and Pan, J. Z. 2021. Geometric models for (temporally) attributed description logics. In *Proceedings of DL*.

d'Amato, C.; Quatraro, N. F.; and Fanizzi, N. 2021. Injecting background knowledge into embedding models for predictive tasks on knowledge graphs. In *Proceedings of ESWC*.

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 *Proceedings of KR*.

Hamilton, W. L.; Bajaj, P.; Zitnik, M.; Jurafsky, D.; and Leskovec, J. 2018. Embedding logical queries on knowledge graphs. In *Proceedings of NeurIPS*.

Helly, E. 1923. Über mengen konvexer körper mit gemeinschaftlichen punkte. *Jahresbericht der Deutschen Mathematiker-Vereinigung* 32:175–176.

Imenes, A.; Guimarães, R.; and Ozaki, A. 2023. Marrying query rewriting and knowledge graph embeddings. In *Proceedings of RuleML+RR*.

Jackermeier, M.; Chen, J.; and Horrocks, I. 2024. Dual box embeddings for the description logic EL^{++}. In *Proceedings of WWW*.

Kazemi, S. M., and Poole, D. 2018. Simple embedding for link prediction in knowledge graphs. In *Proceedings of NeurIPS*.

Kulmanov, M.; Liu-Wei, W.; Yan, Y.; and Hoehndorf, R. 2019. EL embeddings: Geometric construction of models for the description logic EL++. In *Proceedings of IJCAI*.

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).

Mondal, S.; Bhatia, S.; and Mutharaju, R. 2021. EmEL++: Embeddings for EL++ description logic. In *Proceedings of AAAI-MAKE*.

Özçep, Ö. L.; Leemhuis, M.; and Wolter, D. 2020. Cone semantics for logics with negation. In *Proceedings of IJCAI*.

Pavlovic, A., and Sallinger, E. 2023. ExpressivE: A spatio-functional embedding for knowledge graph completion. In *Proceedings of ICLR*.

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

Ren, H., and Leskovec, J. 2020. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In *Proceedings of NeurIPS*.

Ren, H.; Hu, W.; and Leskovec, J. 2020. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In *Proceedings of ICLR*.

Xiong, B.; Potyka, N.; Tran, T.-K.; Nayyeri, M.; and Staab, S. 2022. Faithful embeddings for EL++ knowledge bases. In *Proceedings of ISWC*.

Xiong, B.; Nayyeri, M.; Jin, M.; He, Y.; Cochez, M.; Pan, S.; and Staab, S. 2023. Geometric relational embeddings: A survey. *CoRR* abs/2304.11949.

Zhang, Z.; Wang, J.; Chen, J.; Ji, S.; and Wu, F. 2021. Cone: Cone embeddings for multi-hop reasoning over knowledge graphs. In *Proceedings of NeurIPS*.