



Faithful Embeddings for \mathcal{EL}^{++} Knowledge Bases

Bo Xiong¹(✉), Nico Potyka², Trung-Kien Tran³, Mojtaba Nayyeri¹,
and Steffen Staab^{1,4}

¹ University of Stuttgart, Stuttgart, Germany

{bo.xiong, mojtaba.nayyeri, steffen.staab}@ipvs.uni-stuttgart.de

² Imperial College London, London, UK

n.potyka@imperial.ac.uk

³ Bosch Center for Artificial Intelligence, Renningen, Germany

trungkien.tran@de.bosch.com

⁴ University of Southampton, Southampton, UK

Abstract. Recently, increasing efforts are put into learning continual representations for symbolic knowledge bases (KBs). However, these approaches either only embed the data-level knowledge (ABox) or suffer from inherent limitations when dealing with concept-level knowledge (TBox), i.e., they cannot faithfully model the logical structure present in the KBs. We present BoxEL, a geometric KB embedding approach that allows for better capturing the logical structure (i.e., ABox and TBox axioms) in the description logic \mathcal{EL}^{++} . BoxEL models concepts in a KB as axis-parallel *boxes* that are suitable for modeling concept intersection, entities as points inside boxes, and relations between concepts/entities as *affine transformations*. We show theoretical guarantees (*soundness*) of BoxEL for preserving logical structure. Namely, the learned model of BoxEL embedding with loss 0 is a (logical) model of the KB. Experimental results on (plausible) subsumption reasonings and a real-world application–protein-protein prediction show that BoxEL outperforms traditional knowledge graph embedding methods as well as state-of-the-art \mathcal{EL}^{++} embedding approaches.

Keywords: Ontologies · Knowledge graph embeddings · Semantic web

1 Introduction

Knowledge bases (KBs) provide a *conceptualization* of objects and their relationships, which are of great importance in many applications like biomedical and intelligent systems [5, 25]. KBs are often expressed using description logics (DLs) [3], a family of languages allowing for expressing domain knowledge via logical statements (a.k.a axioms). These logical statements are divided into two parts: 1) an ABox consisting of *assertions* over instances, i.e., factual statements like `isFatherOf(John, Peter)`; 2) a TBox consisting of logical statements constraining concepts, e.g., `Parent \sqsubseteq Person`.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-19433-7_2.

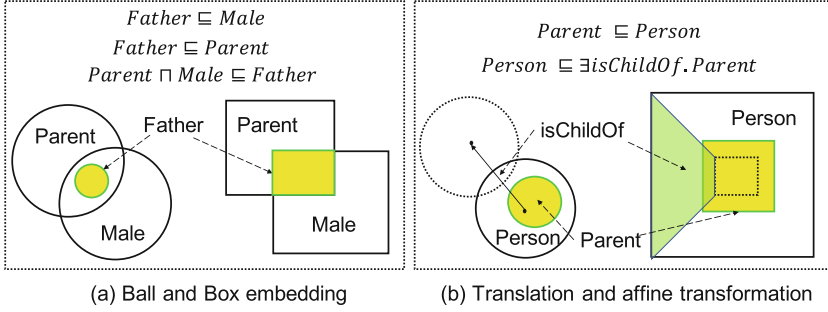


Fig. 1. Two counterexamples of ball embedding and its relational transformation. (a) Ball embedding cannot express concept equivalence $\text{Parent} \cap \text{Male} \equiv \text{Father}$ with intersection operator. (b) The *translation* cannot model relation (e.g. *isChildOf*) between **Person** and **Parent** when they should have different volumes. These two issues can be solved by *box* embedding and modelling relation as *affine transformation* among boxes, respectively.

KBs not only provide clear semantics in the application domains but also enable (classic) reasoners [13, 30] to perform logical inference, i.e., making implicit knowledge explicit. Existing reasoners are highly optimized and scalable but they are limited to only computing classical logical entailment but not designed to perform inductive (analogical) reasoning and cannot handle noisy data. Embedding based methods, which map the objects in the KBs into a low dimensional vector space while keeping the similarity, have been proposed to complement the classical reasoners and shown remarkable empirical performances on performing (non-classical) analogical reasonings.

Most KB embeddings methods [33] focus on embedding data-level knowledge in ABoxes, a.k.a., knowledge graph embeddings (KGEs). However, KGEs cannot preserve concept-level knowledge expressed in TBoxes. Recently, embedding methods for KBs expressed in DLs have been explored. Prominent examples include \mathcal{EL}^{++} [15] that supports conjunction and full existential quantification, and \mathcal{ALC} [23] that further supports logical negation. We focus on \mathcal{EL}^{++} , an underlying formalism of the OWL2 EL profile of the Web Ontology Language [10], which has been used in expressing various biomedical ontologies [5, 25]. For embedding \mathcal{EL}^{++} KBs, several approaches such as Onto2Vec [28] and OPA2Vec [29] have been proposed. These approaches require annotation data and do not model logical structure explicitly. Geometric representations, in which the objects are associated with geometric objects such as balls [15] and convex cones [23], provide a high expressiveness on embedding logical properties. For \mathcal{EL}^{++} KBs, ELEm [15] represents concepts as open n -balls and relations as simple translations. Although effective, ELEm suffers from several major limitations:

- Balls are not closed under intersection and cannot faithfully represent concept intersections. For example, the intersection of two concepts $\text{Parent} \cap \text{Male}$, that is supposed to represent **Father**, is not a ball (see Fig. 1 (a)). Therefore, the concept equivalence $\text{Parent} \cap \text{Male} \equiv \text{Father}$ cannot be captured in the embedding space.
- The relational embedding with simple *translation* causes issues for embedding concepts with varying sizes. For example, Fig. 1 (b) illustrates the embeddings of the axiom $\exists \text{isChildOf}.\text{Person} \sqsubseteq \text{Parent}$ assuming the existence of another axiom

Parent \sqsubseteq Person. In this case, it is impossible to *translate* the larger concept Person into the smaller one Parent,¹ as it does not allow for scaling the size.

- ELEm does not distinguish between entities in ABox and concepts in TBox, but rather regards ABox axioms as special cases of TBox axioms. This simplification cannot fully express the logical structure, e.g., an entity must have minimal volume.

To overcome these limitations, we consider modeling concepts in the KB as *boxes* (i.e., axis-aligned hyperrectangles), encoding entities as points inside the boxes that they should belong to, and the relations as the *affine transformation* between boxes and/or points. Figure 1 (a) shows that the box embedding has closed form of intersection and the *affine transformation* (Fig. 1 (b)) can naturally capture the cases that are not possible in ELEm. In this way, we present BoxEL for embedding \mathcal{EL}^{++} KBs, in which the interpretation functions of \mathcal{EL}^{++} theories in the KB can be represented by the geometric transformations between boxes/points. We formulate BoxEL as an optimization task by designing and minimizing various loss terms defined for each logical statement in the KB. We show theoretical guarantee (*soundness*) of BoxEL in the sense that if the loss of BoxEL embedding is 0, then the trained model is a (logical) model of the KB. Experiments on (plausible) subsumption reasoning over three ontologies and predicting protein-protein interactions show that BoxEL outperforms previous approaches.

2 Related Work

Knowledge graph embeddings (KGEs) have been developed for different tasks. Early works, which focus on link prediction, embed both entities and relations as vectors in a vector space to model the relationships between entities [4, 7, 31]. Prominent examples include *additive* (or *translational*) family [4, 18, 34] and *multiplicative* (or *bilinear*) family [19, 22, 35]. Such techniques only embed the data-level part of KBs and work relatively well for the link prediction tasks. However, KGEs demonstrate limitations when being used to learn the representation of background knowledge such as ontologies of logical rules [9, 23], as well as complex logical query [26, 27].

Inspired by the theory of conceptual spaces [8], several methods have been proposed to embed concepts as convex regions in vector spaces [11], including balls [15] and convex cones [23]. Such conceptual/geometric methods nicely model the set-theoretic semantics that can be used to capture logical rules of knowledge graphs [1], transitive closure in graphs [32] and logical query for multi-hop question answering [26].

Among embeddings for complex concept descriptions, boxes have some conceptual advantages, but they have not been exploited for representing ontologies yet. BoxE [1] does embed some logical rules but mostly focus on embedding the relational patterns in ABoxes. In contrast, our approach BoxEL focuses on \mathcal{EL}^{++} that has a larger TBox and provides soundness guarantees. BoxEL is closely related to ELEm [15], but instead of using ball embedding and translation, we consider box embedding and affine transformation that have various inherent advantages as discussed before. Another difference of

¹ Under the *translation* setting, the embeddings will simply become Parent \equiv Person, which is obviously not what we want as we can express Parent \neq Person with \mathcal{EL}^{++} by propositions like Children \sqcap Parent $\sqsubseteq \perp$, Children \sqsubseteq Person and Children(*a*).

our method is that we use a different encoding that distinguishes between entities and concepts. Furthermore, we take advantage of the volume of boxes for disjointedness representation, resulting in a more natural encoding of the disjointedness of concepts, i.e., two concepts are disjoint iff their intersection has zero volume.

3 Description Logic \mathcal{EL}^{++}

We consider the DL \mathcal{EL}^{++} that underlies multiple biomedical KBs like GALEN [25] and the Gene Ontology [5]. Formally, the syntax of \mathcal{EL}^{++} is built up from a set N_I of *individual names*, N_C of *concept names* and N_R of *role names* (also called *relations*) using the constructors shown in Table 1, where N_I , N_C and N_R are pairwise disjoint.

The semantics of \mathcal{EL}^{++} is defined by *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where the domain $\Delta^{\mathcal{I}}$ is a non-empty set and $\cdot^{\mathcal{I}}$ is a mapping that associates every individual with an element in $\Delta^{\mathcal{I}}$, every concept name with a subset of $\Delta^{\mathcal{I}}$, and every relation name with a relation over $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. An *interpretation* is satisfied if it satisfies the corresponding semantic conditions. The syntax and the corresponding semantics (i.e., interpretation of concept expressions) of \mathcal{EL}^{++} are summarized in Table 1.

An \mathcal{EL}^{++} KB $(\mathcal{A}, \mathcal{T})$ consists of an ABox \mathcal{A} and a TBox \mathcal{T} . The ABox is a set of *concept assertions* $C(a)$ and *role assertions* $r(a, b)$, where C is a concept, r is a relation, and a, b are individuals. The TBox is a set of *concept inclusions* of the form $C \sqsubseteq D$. Intuitively, the ABox contains instance-level information (e.g. $\text{Person}(\text{John})$, $\text{isFatherOf}(\text{John}, \text{Peter})$), while the TBox contains information about concepts (e.g. $\text{Parent} \sqsubseteq \text{Person}$). Every \mathcal{EL}^{++} KB can be transformed such that every TBox statement has the form $C_1 \sqsubseteq D$, $C_1 \sqcap C_2 \sqsubseteq D$, $C_1 \sqsubseteq \exists r.C_2$, $\exists r.C_1 \sqsubseteq D$, where C_1, C_2, D can be the top concept, concept names or nominals and D can also be the bottom concept [2]. The normalized KB can be computed in linear time by introducing new concept names for complex concept expressions and is a conservative extension of the original KB, i.e., every model of the normalized KB is a model of the original KB and every model of the original KB can be extended to be a model of the normalized KB [2].

4 BoxEL for Embedding \mathcal{EL}^{++} Knowledge Bases

In this section, we first present the geometric construction process of \mathcal{EL}^{++} with box embedding and affine transformation, followed by a discussion of the geometric interpretation. Afterward, we describe the BoxEL embedding by introducing proper loss function for each ABox and TBox axiom. Finally, an optimization method is described for the training of BoxEL.

Table 1. Syntax and semantic of \mathcal{EL}^{++} (role inclusions and concrete domains are omitted).

| | Name | Syntax | Semantics |
|--------------|-------------------------|-------------------|--|
| Constructors | Top concept | \top | $\Delta^{\mathcal{I}}$ |
| | Bottom concept | \perp | \emptyset |
| | Nominal | $\{a\}$ | $\{a^{\mathcal{I}}\}$ |
| | Conjunction | $C \sqcap D$ | $C^{\mathcal{I}} \cap D^{\mathcal{I}}$ |
| | Existential restriction | $\exists r.C$ | $\{x \in \Delta^{\mathcal{I}} \mid \exists y \in \Delta^{\mathcal{I}} (x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\}$ |
| ABox | Concept assertion | $C(a)$ | $a^{\mathcal{I}} \in C^{\mathcal{I}}$ |
| | Role assertion | $r(a, b)$ | $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$ |
| TBox | Concept inclusion | $C \sqsubseteq D$ | $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ |

4.1 Geometric Construction

We consider a KB $(\mathcal{A}, \mathcal{T})$ consisting of an ABox \mathcal{A} and a TBox \mathcal{T} where \mathcal{T} has been normalized as explained before. Our goal is to associate entities (or individuals) with points and concepts with boxes in \mathbb{R}^n such that the axioms in the KB are respected.

To this end, we consider two functions m_w, M_w parameterized by a parameter vector w that has to be learned. Conceptually, we consider points as boxes of volume 0. This will be helpful later to encode the meaning of axioms for points and boxes in a uniform way. Intuitively, $m_w : N_I \cup N_C \rightarrow \mathbb{R}^n$ maps individual and concept names to the lower left corner and $M_w : N_I \cup N_C \rightarrow \mathbb{R}^n$ maps them to the upper right corner of the box that represents them. For individuals $a \in N_I$, we have $m_w(a) = M_w(a)$, so that it is sufficient to store only one of them. The *box associated with C* is defined as

$$\text{Box}_w(C) = \{x \in \mathbb{R}^n \mid m_w(C) \leq x \leq M_w(C)\}, \quad (1)$$

where the inequality is defined component-wise.

Note that boxes are closed under intersection, which allows us to compute the volume of the intersection of boxes. The lower corner of the box $\text{Box}_w(C) \cap \text{Box}_w(D)$ is $\max(m_w(C), m_w(D))$ and the upper corner is $\min(M_w(C), M_w(D))$, where minimum and maximum are taken component-wise. The volume of boxes can be used to encode axioms in a very concise way. However, as we will describe later, one problem is that points have volume 0. This does not allow distinguishing empty boxes from points. To show that our encoding correctly captures the logical meaning of axioms, we will consider a *modified volume* that assigns a non-zero volume to points and some empty boxes. The (modified) volume of a box is defined as

$$\text{MVol}(\text{Box}_w(C)) = \prod_{i=1}^n \max(0, M_w(C)_i - m_w(C)_i + \epsilon), \quad (2)$$

where $\epsilon > 0$ is a small constant. A point now has volume ϵ^n . Some empty boxes can actually have arbitrarily large modified volume. For example the 2D-box with lower corner $(0, 0)$ and upper corner $(-\frac{\epsilon}{2}, N)$ has volume $\frac{\epsilon \cdot N}{2}$. While this is not meaningful geometrically, it does not cause any problems for our encoding because we only want to ensure that boxes with zero volume are empty (and not points). In practice, we will use *softplus volume* as approximation (see Sect. 4.5).

We associate every role name $r \in N_r$ with an affine transformation denoted by $T_w^r(x) = D_w^r x + b_w^r$, where D_w^r is an $(n \times n)$ diagonal matrix with non-negative entries and $b_w^r \in \mathbb{R}^n$ is a vector. In a special case where all diagonal entries of D_w^r are -1, $T_w^r(x)$ captures translations. Note that relations have been represented by translation vectors analogous to TransE in [15]. However, this necessarily means that the concept associated with the range of a role has the same size as its domain. This does not seem very intuitive, in particular, for N-to-one relationships like *has_nationality* or *lives_in* that map many objects to the same object. Note that $T_w^r(\text{Box}_w(C)) = \{T_w^r(x) \mid x \in \text{Box}_w(C)\}$ is the box with lower corner $T_w^r(m_w(C))$ and upper corner $T_w^r(M_w(C))$. To show this, note that $m_w(C) < M_w(C)$ implies $D_w^r m_w(C) \leq D_w^r M_w(C)$ because D_w^r is a diagonal matrix with non-negative entries. Hence, $T_w^r(m_w(C)) = D_w^r m_w(C) + b_w^r \leq D_w^r M_w(C) + b_w^r = T_w^r(M_w(C))$. For $m_w(C) \geq M_w(C)$, both $\text{Box}_w(C)$ and $T_w^r(\text{Box}_w(C))$ are empty.

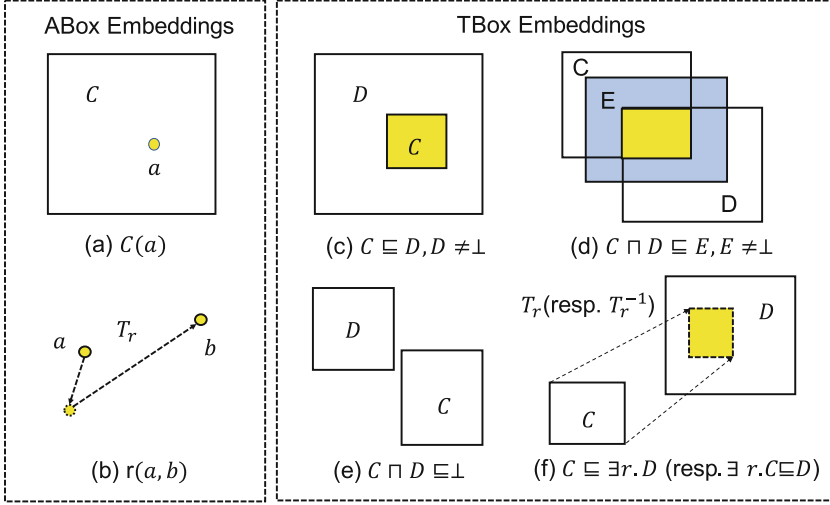


Fig. 2. The geometric interpretation of logical statements in ABox (*left*) and TBox (*right*) expressed by DL \mathcal{EL}^{++} with BoxEL embeddings. The concepts are represented by *boxes*, entities are represented by *points* and relations are represented by *affine transformations*. T_r and T_r^{-1} denote the transformation function of relation r and its inverse function, respectively.

Overall, we have the following parameters:

- for every individual name $a \in N_I$, we have n parameters for the vector $m_w(a)$ (since $m_w(a) = M_w(A)$, we have to store only one of m_w and M_w),
- for every concept name $C \in N_C$, we have $2n$ parameters for the vectors $m_w(C)$ and $M_w(C)$,
- for every role name $r \in N_r$, we have $2n$ parameters. n parameters for the diagonal elements of D_w^r and n parameters for the components of b_w^r .

As we explained informally before, w summarizes all parameters. The overall number of parameters in w is $n \cdot (|N_I| + 2 \cdot |N_C| + 2 \cdot |N_r|)$.

4.2 Geometric Interpretation

The next step is to encode the axioms in our KB. However, we do not want to do this in an arbitrary fashion, but, ideally, in a way that gives us some analytical guarantees. [15] made an interesting first step by showing that their encoding is *sound*. In order to understand soundness, it is important to know that the parameters of the embedding are learnt by minimizing a loss function that contains a loss term for every axiom. Soundness then means that if the loss function yields 0, then the KB is satisfiable. Recall that satisfiability means that there is an interpretation that satisfies all axioms in the KB. Ideally, we should be able to construct such an interpretation directly from our embedding. This is indeed what the authors in [15] did. The idea is that points in the vector space make up the domain of the interpretation, the points that lie in

regions associated with concepts correspond to the interpretation of this concept and the interpretation of roles correspond to translations between points like in TransE. In our context, geometric interpretation can be defined as follows.

Definition 1 (Geometric Interpretation). *Given a parameter vector w representing an \mathcal{EL}^{++} embedding, the corresponding geometric interpretation $\mathcal{I}_w = (\Delta^{\mathcal{I}_w}, \cdot^{\mathcal{I}_w})$ is defined as follows:*

1. $\Delta^{\mathcal{I}_w} = \mathbb{R}^n$,
2. for every concept name $C \in N_C$, $C^{\mathcal{I}_w} = \text{Box}_w(C)$,
3. for every role $r \in N_R$, $r^{\mathcal{I}_w} = \{(x, y) \in \Delta^{\mathcal{I}_w} \times \Delta^{\mathcal{I}_w} \mid T_w^r(x) = y\}$,
4. for every individual name $a \in N_I$, $a^{\mathcal{I}_w} = m_w(a)$.

We will now encode the axioms by designing one loss term for every axiom in a normalized \mathcal{EL}^{++} KB, such that the axiom is satisfied by the geometric interpretation when the loss is 0. All proofs of propositions are attached in the supplementary material.

4.3 ABox Embedding

ABox contains concept assertions and role assertions. We introduce the following two loss terms that respect the geometric interpretations.

Concept Assertion. Geometrically, a concept assertion $C(a)$ asserts that the point $m_w(a)$ is inside the box $\text{Box}_w(C)$ (see Fig. 2 (a)). This can be expressed by demanding $m_w(C) \leq m_w(a) \leq M_w(C)$ for every component. The loss $\mathcal{L}_{C(a)}(w)$ is defined by

$$\mathcal{L}_{C(a)}(w) = \sum_{i=1}^n \|\max(0, m_w(a)_i - M_w(C)_i)\|_2 + \sum_{i=1}^n \|\max(0, m_w(C)_i - m_w(a)_i)\|_2.$$

Role Assertion. Geometrically, a role assertion $r(a, b)$ means that the point $m_w(a)$ should be mapped to $m_w(b)$ by the transformation T_w^r (see Fig. 2 (b)). That is, we should have $T_w^r(m_w(a)) = m_w(b)$. We define a loss term

$$\mathcal{L}_{r(a,b)}(w) = \|T_w^r(m_w(a)) - m_w(b)\|_2. \quad (3)$$

It is clear from the definition that when the loss terms are 0, the axioms are satisfied in their geometric interpretation.

Proposition 1. *We have*

1. If $\mathcal{L}_{C(a)}(w) = 0$, then $\mathcal{I}_w \models C(a)$,
2. If $\mathcal{L}_{r(a,b)}(w) = 0$, then $\mathcal{I}_w \models r(a, b)$.

4.4 TBox Embedding

For the TBox, we define loss terms for the four cases in the normalized KB. Before doing so, we define an auxiliary function that will be used inside these loss terms.

Definition 2 (Disjoint measurement). Given two boxes B_1, B_2 , the disjoint measurement can be defined by the (modified) volumes of B_1 and the intersection box $B_1 \cap B_2$,

$$\text{Disjoint}(B_1, B_2) = 1 - \frac{\text{MVol}(B_1 \cap B_2)}{\text{MVol}(B_1)}. \quad (4)$$

We have the following guarantees.

- Lemma 1.** 1. $0 \leq \text{Disjoint}(B_1, B_2) \leq 1$,
 2. $\text{Disjoint}(B_1, B_2) = 0$ implies $B_1 \subseteq B_2$,
 3. $\text{Disjoint}(B_1, B_2) = 1$ implies $B_1 \cap B_2 = \emptyset$.

NF1: Atomic Subsumption. An axiom of the form $C \sqsubseteq D$ geometrically means that $\text{Box}_w(C) \subseteq \text{Box}_w(D)$ (see Fig. 2 (c)). If $D \neq \perp$, we consider the loss term

$$\mathcal{L}_{C \sqsubseteq D}(w) = \text{Disjoint}(\text{Box}_w(C), \text{Box}_w(D)). \quad (5)$$

For the case $D = \perp$ where C is not a nominal, e.g., $C \sqsubseteq \perp$, we define the loss term

$$\mathcal{L}_{C \sqsubseteq \perp}(w) = \max(0, M_w(C)_0 - m_w(C)_0 + \epsilon). \quad (6)$$

If C is a nominal, the axiom is inconsistent and our model can just return an error.

Proposition 2. If $\mathcal{L}_{C \sqsubseteq D}(w) = 0$, then $\mathcal{I}_w \models C \sqsubseteq D$, where we exclude the inconsistent case $C = \{a\}, D = \perp$.

NF2: Conjunct Subsumption. An axiom of the form $C \sqcap D \sqsubseteq E$ means that $\text{Box}(C) \cap \text{Box}(D) \subseteq \text{Box}(E)$ (see Fig. 2 (d)). Since $\text{Box}(C) \cap \text{Box}(D)$ is a box again, we can use the same idea as for NF1. For the case $E \neq \perp$, we define the loss term as

$$\mathcal{L}_{C \sqcap D \sqsubseteq E}(w) = \text{Disjoint}(\text{Box}_w(C) \cap \text{Box}_w(D), \text{Box}_w(E)). \quad (7)$$

For $E = \perp$, the axiom states that C and D must be disjoint. The disjointedness can be interpreted as the volume of the intersection of the associated boxes being 0 (see Fig. 2 (e)). However, just using the volume as a loss term may not work well because a minimization algorithm may minimize the volume of the boxes instead of the volume of their intersections. Therefore, we normalize the loss term by dividing by the volume of the boxes. Given by

$$\mathcal{L}_{C \sqcap D \sqsubseteq \perp}(w) = \frac{\text{MVol}(\text{Box}_w(C) \cap \text{Box}_w(D))}{\text{MVol}(\text{Box}_w(C)) + \text{MVol}(\text{Box}_w(D))}. \quad (8)$$

Proposition 3. If $\mathcal{L}_{C \sqcap D \sqsubseteq E}(w) = 0$, then $\mathcal{I}_w \models C \sqcap D \sqsubseteq E$, where we exclude the inconsistent case $a \sqcap a \sqsubseteq \perp$ (that is, $C = D = \{a\}, E = \perp$).

NF3: Right Existential. Next, we consider axioms of the form $C \sqsubseteq \exists r.D$. Note that $\exists r.D$ describes those entities that are in relation r with an entity from D . Geometrically, those are points that are mapped to points in $\text{Box}_w(D)$ by the affine transformation

corresponding to r . $C \sqsubseteq \exists r.D$ then means that every point in $\text{Box}_w(C)$ must be mapped to a point in $\text{Box}_w(D)$, that is the mapping of $\text{Box}_w(C)$ is contained in $\text{Box}_w(D)$ (see Fig. 2 (f)). Therefore, the encoding comes again down to encoding a subset relationship as before. The only difference to the first normal form is that $\text{Box}_w(C)$ must be mapped by the affine transformation T_w^r . These considerations lead to the following loss term

$$\mathcal{L}_{C \sqsubseteq \exists r.D}(w) = \text{Disjoint}(T_w^r(\text{Box}_w(C)), \text{Box}_w(D)). \quad (9)$$

Proposition 4. *If $\mathcal{L}_{C \sqsubseteq \exists r.D}(w) = 0$, then $\mathcal{I}_w \models C \sqsubseteq \exists r.D$.*

NF4: Left Existential. Axioms of the form $\exists r.C \sqsubseteq D$ can be treated symmetrically to the previous case (see Fig. 2 (f)). We only consider the case $D \neq \perp$ and define the loss

$$\mathcal{L}_{\exists r.C \sqsubseteq D}(w) = \text{Disjoint}(T_w^{-r}(\text{Box}_w(C)), \text{Box}_w(D)), \quad (10)$$

where T_w^{-r} is the inverse function of T_w^r that is defined by $T_w^{-r}(x) = D_w^{-r}x - D_w^{-r}b_w^r$, where D_w^{-r} is obtained from D_w^r by replacing all diagonal elements with their reciprocal. Strictly speaking, the inverse only exists if all diagonal entries of D_w^r are non-zero. However, we assume that the entries that occur in a loss term of the form $\mathcal{L}_{\exists r.C \sqsubseteq D}(w)$ remain non-zero in practice when we learn them iteratively.

Proposition 5. *If $\mathcal{L}_{\exists r.C \sqsubseteq D}(w) = 0$, then $\mathcal{I}_w \models \exists r.C \sqsubseteq D$.*

4.5 Optimization

Softplus Approximation. For optimization, while the computation of the volume of boxes is straightforward, using a precise *hard volume* is known to cause problems when learning the parameters using gradient descent algorithms, e.g. there is no training signal (gradient flow) when box embeddings that should overlap but become disjoint [6, 17, 24]. To mitigate the problem, we approximate the volume of boxes by the *softplus volume* [24] due to its simplicity.

$$\text{SVol}(\text{Box}_w(C)) = \prod_{i=1}^n \text{Softplus}_t(M_w(C)_i - m_w(C)_i) \quad (11)$$

where t is a temperature parameter. The softplus function is defined as $\text{softplus}_t(x) = t \log(1 + e^{x/t})$, which can be regarded as a smoothed version of the ReLU function ($\max\{0, x\}$) used for calculating the volume of *hard boxes*. In practice, the *softplus volume* is used to replace the *modified volume* in Eq. (2) as it empirically resolves the same issue that point has zero volume.

Regularization. We add a regularization term in Eq. (12) to all non-empty boxes to encourage that the boxes lie in the unit box $[0, 1]^n$.

$$\lambda = \sum_{i=1}^n \max(0, M_w(C)_i - 1 + \epsilon) + \max(0, -m_w(C)_i - \epsilon) \quad (12)$$

In practice, this also avoids numerical stability issues. For example, to minimize a loss term, a box that should have a fixed volume could become very *slim*, i.e. some side lengths be extremely large while others become extremely small.

Negative Sampling. In principle, the embeddings can be optimized without negatives. However, we empirically find that the embeddings will be highly overlapped without negative sampling. e.g. for role assertion $r(a, b)$, a and b will simply become the same point. We generate negative samples for the role assertion $r(a, b)$ by randomly replacing one of the head or tail entity. Finally, we sum up all the loss terms, and learn the embeddings by minimizing the loss with Adam optimizer [14].

5 Empirical Evaluation

5.1 A Proof-of-Concept Example

We begin by first validating the model in modeling a toy ontology–family domain [15], which is described by the following axioms:²

| | |
|---|--|
| Male \sqsubseteq Person | Female \sqsubseteq Person |
| Father \sqsubseteq Male | Mother \sqsubseteq Female |
| Father \sqsubseteq Parent | Mother \sqsubseteq Parent |
| Female \sqcap Male $\sqsubseteq \perp$ | Female \sqcap Parent \sqsubseteq Mother |
| Male \sqcap Parent \sqsubseteq Father | \exists hasChild.Person \sqsubseteq Parent |
| Parent \sqsubseteq Person | Parent $\sqsubseteq \exists$ hasChild.Person |
| Father(Alex) | Father(Bob) |
| Mother(Marie) | Mother(Alice) |

We set the dimension to 2 to visualize the embeddings. Figure 3 shows that the generated embeddings accurately encode all of the axioms. In particular, the embeddings of **Father** and **Mother** align well with the conjunction **Parent \sqcap Male** and **Parent \sqcap Female**, respectively, which is impossible to be achieved by ELEM.

5.2 Subsumption Reasoning

We evaluate the effectiveness of BoxEL on (plausible) subsumption reasoning (also known as ontology completion). The problem is to predict whether a concept is subsumed by another one. For each subsumption pair $C \sqsubseteq D$, the scoring function can be defined by

$$P(C \sqsubseteq D) = \frac{\text{MVOL}(\text{Box}(C) \cap \text{Box}(D))}{\text{MVOL}(\text{Box}(C))}. \quad (13)$$

Note that such subsumption relations are not necessary to be (logically) entailed by the input KB, e.g., a subsumption relation can be plausibly inferred by $P(C \sqsubseteq D) =$

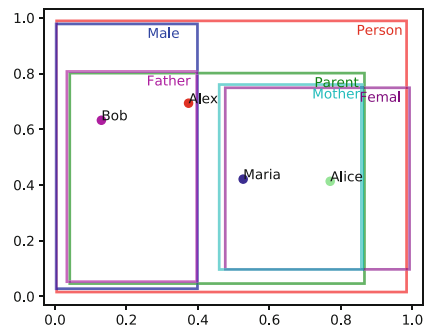


Fig. 3. BoxEL embeddings in the family domain.

² Compared with the example given in [15], we add additional concept assertion statements that distinguish entities and concepts.

Table 2. Summary of classes, relations and axioms in different ontologies. NF_i represents the i^{th} normal form.

| Ontology | GO | GALEN | ANATOMY |
|-----------|-------|-------|---------|
| Classes | 45895 | 24353 | 106363 |
| Relations | 9 | 1010 | 157 |
| NF1 | 85480 | 28890 | 122142 |
| NF2 | 12131 | 13595 | 2121 |
| NF3 | 20324 | 28118 | 152289 |
| NF4 | 12129 | 13597 | 2143 |
| Disjoint | 30 | 0 | 184 |

Table 3. The accuracies (for which the prediction is true if and only if the subclass box is exactly inside the superclass box) achieved by the embeddings of different approaches in terms of geometric interpretation of the classes in various ontologies.

| | ELEm | EmEL ⁺⁺ | BoxEL |
|---------|-------|--------------------|--------------|
| GO | 0.250 | 0.415 | 0.489 |
| GALEN | 0.480 | 0.345 | 0.788 |
| ANATOMY | 0.069 | 0.215 | 0.453 |

Table 4. The ranking based measures of embedding models for subsumption reasoning on the testing set. * denotes the results from [20].

| Dataset | Metric | TransE* | TransH* | DistMult* | ELEm | EmEL ⁺⁺ | BoxEL |
|---------|-----------|---------|---------|-----------|-------|--------------------|-------|
| GO | Hits@10 | 0.00 | 0.00 | 0.00 | 0.09 | 0.10 | 0.03 |
| | Hits@100 | 0.00 | 0.00 | 0.00 | 0.16 | 0.22 | 0.08 |
| | AUC | 0.53 | 0.44 | 0.50 | 0.70 | 0.76 | 0.81 |
| | Mean Rank | - | - | - | 13719 | 11050 | 8980 |
| GALEN | Hits@10 | 0.00 | 0.00 | 0.00 | 0.07 | 0.10 | 0.02 |
| | Hits@100 | 0.00 | 0.00 | 0.00 | 0.14 | 0.17 | 0.03 |
| | AUC | 0.54 | 0.48 | 0.51 | 0.64 | 0.65 | 0.85 |
| | Mean Rank | - | - | - | 8321 | 8407 | 3584 |
| ANATOMY | Hits@10 | 0.00 | 0.00 | 0.00 | 0.18 | 0.18 | 0.03 |
| | Hits@100 | 0.01 | 0.00 | 0.00 | 0.38 | 0.40 | 0.04 |
| | AUC | 0.53 | 0.44 | 0.49 | 0.73 | 0.76 | 0.91 |
| | Mean Rank | - | - | - | 28564 | 24421 | 10266 |

0.9, allowing for non-classical plausible reasoning. While the subsumption reasoning does not need negatives, we add an additional regularization term for non-subsumption axiom. In particular, for each atomic subsumption axiom $C \sqsubseteq D$, we generate a non-subsumption axiom $C \not\sqsubseteq D'$ or $C' \not\sqsubseteq D$ by randomly replacing one of the concepts C and D . Note that this does not produce regular negative samples as the generated concepts pair does not have to be disjoint. Thus, the loss term for non-subsumption axiom cannot be simply defined by $\mathcal{L}_{C \not\sqsubseteq D'} = 1 - \mathcal{L}_{C \sqsubseteq D}$. Instead, we define the loss term as $\mathcal{L}_{C \not\sqsubseteq D'} = \phi(1 - \mathcal{L}_{C \sqsubseteq D})$ by multiplying a small positive constant ϕ that encourages splitting the non-subsumption concepts while does not encourage them to be disjoint. If $\phi = 1$, the loss would encourage the non-subsumption concepts to be disjoint. We empirically show that $\phi = 1$ produces worse performance as we do not want non-subsumption concepts to be disjoint.

Datasets. We use three biomedical ontologies as our benchmark. 1) **Gene Ontology (GO)** [12] integrates the representation of genes and their functions across all species. 2) **GALEN** [25] is a clinical ontology. 3) **Anatomy** [21] is an ontology that represents linkages of different phenotypes to genes. Table 2 summarizes the statistical information of these datasets. The subclass relations are split into training set (70%), validation set (20%) and testing set (10%), respectively.

Evaluation Protocol. Two strategies can be used to measure the effectiveness of the embeddings. 1) Ranking based measures rank the probability of C subsumed by all concepts. We evaluate and report four ranking based measures. Hits@10, Hits@100 describe the fraction of true cases that appear in the first 10 and 100 test cases of the sorted rank list, respectively. Mean rank computes the arithmetic mean over all individual ranks (i.e. $MR = \frac{1}{|\mathcal{I}|} \sum_{rank \in \mathcal{I}} rank$, where $rank$ is the individual rank), while AUC computes the area under the ROC curve. 2) Accuracy based measure is a stricter criterion, for which the prediction is true if and only if the subclass box is exactly inside the superclass box (even not allowing the subclass box slightly outside the superclass box). We use this measure as it evaluates the performance of embeddings on retaining the underlying characteristics of ontology in vector space. We only compare ELEM and EmEL⁺⁺ as KGE baselines fail in this setting (KGEs cannot preserve the ontology).

Implementation Details. The ontology is normalized into standard normal forms, which comprise a set of axioms that can be used as the *positive samples*. Similar to previous works [15], we perform normalization using the OWL APIs and the APIs offered by the jCel reasoner [18]. The hyperparameter for negative sampling is set to $\phi = 0.05$. For ELEM and EmEL⁺⁺, the embedding size is searched from $n = [50, 100, 200]$ and margin parameter is searched from $\gamma = [-0.1, 0, 0.1]$. Since box embedding has double the number of parameters of ELEM and EmEL⁺⁺, we search the embedding size from $n = [25, 50, 100]$ for BoxEL. We summarize the best performing hyperparameters in our supplemental material. All experiments are evaluated with 10 random seeds and the mean results are reported for numerical comparisons.

Baselines. We compare the state-of-the-art \mathcal{EL}^{++} embeddings (ELEM) [15], the first geometric embeddings of \mathcal{EL}^{++} , as well as the extension EmEL⁺⁺ [20] that additionally considers the role inclusion and role chain embedding, as our major baselines. For comparison with classical methods, we also include the reported results of three classical KGEs in [20], including TransE [4], TransH [34] and DistMult [35].

Results. Table 4 summarizes the ranking based measures of embedding models. We first observe that both ELEM and EmEL⁺⁺ perform much better than the three standard KGEs (TransE, TransH, and DistMult) on all three datasets, especially on hits@k for which KGEs fail, showcasing the limitation of KGEs and the benefits of geometric embeddings on encoding logic structures. EmEL⁺⁺ performs slightly better than ELEM on all three datasets. Overall, our model BoxEL outperforms ELEM and EmEL⁺⁺. In

particular, we find that for Mean Rank and AUC, our model achieves significant performance gains on all three datasets. Note that Mean Rank and AUC have theoretical advantages over hits@k because hits@k is sensitive to any model performance changes while Mean Rank and AUC reflect the average performance, demonstrating that BoxEL achieves better average performance. Table 4 shows the accuracies of different embeddings in terms of the geometric interpretation of the classes in various ontologies. It clearly demonstrates that BoxEL outperforms ELEM and EmEL⁺⁺ by a large margin, showcasing that BoxEL preserves the underlying ontology characteristics in vector space better than ELEM and EmEL⁺⁺ that use ball embeddings.

Table 5. Prediction performance on protein-protein interaction (yeast).

| Method | Raw hits@10 | Filtered hits@10 | Raw hits@100 | Filtered hits@100 | Raw mean Rank | Filtered mean rank | Raw AUC | Filtered AUC |
|--------------------|-------------|------------------|--------------|-------------------|---------------|--------------------|-------------|--------------|
| TransE | 0.06 | 0.13 | 0.32 | 0.40 | 1125 | 1075 | 0.82 | 0.83 |
| BoxE | 0.08 | 0.14 | 0.36 | 0.43 | 633 | 620 | 0.85 | 0.85 |
| SimResnik | 0.09 | 0.17 | 0.38 | 0.48 | 758 | 707 | 0.86 | 0.87 |
| SimLin | 0.08 | 0.15 | 0.33 | 0.41 | 875 | 825 | 0.8 | 0.85 |
| ELEM | 0.08 | 0.17 | 0.44 | 0.62 | 451 | 394 | 0.92 | 0.93 |
| EmEL ⁺⁺ | 0.08 | 0.16 | 0.45 | 0.63 | 451 | 397 | 0.90 | 0.91 |
| Onto2Vec | 0.08 | 0.15 | 0.35 | 0.48 | 641 | 588 | 0.79 | 0.80 |
| OPA2Vec | 0.06 | 0.13 | 0.39 | 0.58 | 523 | 467 | 0.87 | 0.88 |
| BoxEL | 0.09 | 0.20 | 0.52 | 0.73 | 423 | 379 | 0.93 | 0.94 |

Table 6. Prediction performance on protein-protein interaction (human).

| Method | Raw hits@10 | Filtered hits@10 | Raw hits@100 | Filtered hits@100 | Raw mean rank | Filtered mean rank | Raw AUC | Filtered AUC |
|--------------------|-------------|------------------|--------------|-------------------|---------------|--------------------|-------------|--------------|
| TransE | 0.05 | 0.11 | 0.24 | 0.29 | 3960 | 3891 | 0.78 | 0.79 |
| BoxE | 0.05 | 0.10 | 0.26 | 0.32 | 2121 | 2091 | 0.87 | 0.87 |
| SimResnik | 0.05 | 0.09 | 0.25 | 0.30 | 1934 | 1864 | 0.88 | 0.89 |
| SimLin | 0.04 | 0.08 | 0.20 | 0.23 | 2288 | 2219 | 0.86 | 0.87 |
| ELEM | 0.01 | 0.02 | 0.22 | 0.26 | 1680 | 1638 | 0.90 | 0.90 |
| EmEL ⁺⁺ | 0.01 | 0.03 | 0.23 | 0.26 | 1671 | 1638 | 0.90 | 0.91 |
| Onto2Vec | 0.05 | 0.08 | 0.24 | 0.31 | 2435 | 2391 | 0.77 | 0.77 |
| OPA2Vec | 0.03 | 0.07 | 0.23 | 0.26 | 1810 | 1768 | 0.86 | 0.88 |
| BoxEL (Ours) | 0.07 | 0.10 | 0.42 | 0.63 | 1574 | 1530 | 0.93 | 0.93 |

5.3 Protein-Protein Interactions

Dataset. We use a biomedical knowledge graph built by [15] from Gene Ontology (TBox) and STRING database (ABox) to conduct this task. Gene Ontology contains information about the functions of proteins, while STRING database consists of the protein-protein interactions. We use the protein-protein interaction data of yeast and human organisms, respectively. For each pair of proteins (P_1, P_2) that exists in STRING, we add a role assertion $\text{interacts}(P_1, P_2)$. If protein P is associated with

the function F , we add a membership axiom $\{P\} \sqsubseteq \exists \text{hasFunction}.F$, the membership assertion can be regarded as a special case of NF3, in which P is a point (i.e. zero-volume box). The interaction pairs of proteins are split into training (80%), testing (10%) and validation (10%) sets. To perform prediction for each protein pair (P_1, P_2) , we predict whether the role assertion $\text{interacts}(P_1, P_2)$ hold. This can be measured by Eq. (14).

$$P(\text{interacts}(P_1, P_2)) = \left\| T_w^{\text{interacts}}(m_w(P_1)) - m_w(P_2) \right\|_2. \quad (14)$$

where $T_w^{\text{interacts}}$ is the affine transformation function for relation interacts . For each positive interaction pair $\text{interacts}(P_1, P_2)$, we generate a corrupted negative sample by randomly replacing one of the head and tail proteins.

Baselines. We consider ELEM [15] and EmEL⁺⁺ [20] as our two major baselines as they have been shown outperforming the traditional KGEs. We also report the result of Onto2Vec [28] that treats logical axioms as a text corpus and OPA2Vec [29] that combines logical axioms with annotation properties. Besides, we report the results of two semantic similarity measures: Resnik’s similarity and Lin’s similarity in [15]. For KGEs, we compare TransE [4]) and BoxE [1]. We report the hits@10, hits@100, mean rank and AUC (area under the ROC curve) as explained before for numerical comparison. Both raw ranking measures and filtered ranking measures that ignore the triples that are already known to be true in the training stage are reported. Baseline results are taken from the standard benchmark developed by [16].³

Table 7. The performance of BoxEL with affine transformation (AffineBoxEL) and BoxEL with translation (TransBoxEL) on yeast protein-protein interaction.

| Method | EmEL | | TransBoxEL | | AffineBoxEL | |
|-----------|------|----------|-------------|----------|-------------|-------------|
| | Raw | Filtered | Raw | Filtered | Raw | Filtered |
| Hits@10 | 0.08 | 0.17 | 0.04 | 0.18 | 0.09 | 0.20 |
| Hits@100 | 0.44 | 0.62 | 0.54 | 0.68 | 0.52 | 0.73 |
| Mean rank | 451 | 394 | 445 | 390 | 423 | 379 |
| AUC | 0.92 | 0.93 | 0.93 | 0.93 | 0.93 | 0.94 |

Table 8. The performance of BoxEL with point entity embedding and box entity embedding on yeast protein-protein interaction dataset.

| Method | EmEL | | BoxEL (boxes) | | BoxEL (points) | |
|-----------|------|----------|---------------|----------|----------------|-------------|
| | Raw | Filtered | Raw | Filtered | Raw | Filtered |
| Hits@10 | 0.08 | 0.17 | 0.09 | 0.19 | 0.09 | 0.20 |
| Hits@100 | 0.44 | 0.62 | 0.48 | 0.68 | 0.52 | 0.73 |
| Mean rank | 451 | 394 | 450 | 388 | 423 | 379 |
| AUC | 0.92 | 0.93 | 0.92 | 0.93 | 0.93 | 0.94 |

Overall Results. Table 5 and Table 6 summarize the performance of protein-protein prediction in yeast and human organisms, respectively. We first observe that similarity based methods (SimResnik and SimLin) roughly outperform TransE, showcasing the limitation of classical knowledge graph embeddings. BoxE roughly outperforms TransE as it does encode some logical properties. The geometric methods ELEM and EmEL⁺⁺ fail on the hits@10 measures and does not show significant performance gains on the hits@100 measures in human dataset. However, ELEM and EmEL⁺⁺ outperform TransE, BoxE and similarity based methods on Mean Rank and AUC by a

³ <https://github.com/bio-ontology-research-group/machine-learning-with-ontologies>.

large margin, especially for the Mean Rank, showcasing the expressiveness of geometric embeddings. Onto2Vec and OPA2Vec achieve relatively better results than TransE and similarity based methods, but cannot compete ELEM and EmEL⁺⁺. We conjecture that this is due to the fact that they mostly consider annotation information but cannot encode the logical structure explicitly. Our method, BoxEL consistently outperforms all methods in hits@100, Mean Rank and AUC in both datasets, except the competitive results of hits@10, showcasing the better expressiveness of BoxEL.

5.4 Ablation Studies

Transformation vs Translation. To study the contributions of using boxes for modeling concepts and using affine transformation for modeling relations, we conduct an ablation study by comparing relation embeddings with affine transformation (AffineBoxEL) and translation (TransBoxEL). The only difference of TransBox to the AffineBox is that TransBox does not associate a scaling factor for each relation. Table 7 clearly shows that TransBoxEL outperforms EmEL⁺⁺, showcasing the benefits of box modeling compared with ball modeling. While AffineBoxEL further improves TransBoxEL, demonstrating the advantages of affine transformation. Hence, we could conclude that both of our proposed entity and relation embedding components boost the performance.

Entities as Points vs Boxes. As mentioned before, distinguishing entities and concepts by identifying entities as points has better theoretical properties. Here, we study how this distinction influences the performance. For this purpose, we eliminate the ABox axioms by replacing each individual with a singleton class and rewriting relation assertions $r(a, b)$ and concept assertions $C(a)$ as $\{a\} \sqsubseteq \exists r.\{b\}$ and $\{a\} \sqsubseteq C$, respectively. In this case, we only have TBox embeddings and the entities are embedded as regular boxes. Table 8 shows that for hits@k, there is marginal significant improvement of point entity embedding over boxes entity embedding, however, point entity embedding consistently outperforms box entity embedding on Mean Rank and AUC, showcasing the benefits of distinguishing entities and concepts.

6 Conclusion

This paper proposes BoxEL, a geometric KB embedding method that explicitly models the logical structure expressed by the theories of \mathcal{EL}^{++} . Different from the standard KGEs that simply ignore the analytical guarantees, BoxEL provides *soundness* guarantee for the underlying logical structure by incorporating background knowledge into machine learning tasks, offering a more reliable and logic-preserved fashion for KB reasoning. The empirical results further demonstrate that BoxEL outperforms previous KGEs and \mathcal{EL}^{++} embedding approaches on subsumption reasoning over three ontologies and predicting protein-protein interactions in a real-world biomedical KB.

Supplemental Material Statement: Source code and datasets are available for reproducing the results.⁴ Full proofs of propositions and lemmas are available in a long version of the paper.⁵

Acknowledgments. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Bo Xiong. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No: 860801. Nico Potyka was partially funded by DFG projects Evowipe/COFFEE.

References

1. Abboud, R., Ceylan, İ.İ., Lukasiewicz, T., Salvatori, T.: Boxe: A box embedding model for knowledge base completion. In: *NeurIPS (2020)*
2. Baader, F., Brandt, S., Lutz, C.: Pushing the el envelope. In: *IJCAI*. vol. 5, pp. 364–369 (2005)
3. Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., Nardi, D., et al.: *The description logic handbook: Theory, implementation and applications*. Cambridge University Press (2003)
4. Bordes, A., Usunier, N., García-Durán, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: *NIPS*. pp. 2787–2795 (2013)
5. Consortium, G.O.: Gene ontology consortium: going forward. *Nucleic acids research* **43**(D1), D1049–D1056 (2015)
6. Dasgupta, S.S., Boratko, M., Zhang, D., Vilnis, L., Li, X., McCallum, A.: Improving local identifiability in probabilistic box embeddings. In: *NeurIPS (2020)*
7. Dettmers, T., Minervini, P., Stenetorp, P., Riedel, S.: Convolutional 2d knowledge graph embeddings. In: *AAAI*, pp. 1811–1818. AAAI Press (2018)
8. Gärdenfors, P.: *Conceptual spaces - the geometry of thought*. MIT Press (2000)
9. Garg, D., Ikbāl, S., Srivastava, S.K., Vishwakarma, H., Karanam, H.P., Subramaniam, L.V.: Quantum embedding of knowledge for reasoning. In: *NeurIPS*, pp. 5595–5605 (2019)
10. Graua, B.C., Horrocks, I., Motika, B., Parsiab, B., Patel-Schneider, P., Sattler, U.: Web semantics: science, services and agents on the world wide web. *Web Semantics: Sci. Serv. Agents World Wide Web* **6**, 309–322 (2008)
11. Gutiérrez-Basulto, V., Schockaert, S.: From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules. In: *KR*, pp. 379–388. AAAI Press (2018)
12. Harris, M., et al.: The gene ontology (go) database and informatics resource nucleic acids research, **32**. D258–D261 (2004)
13. Kazakov, Y., Krötzsch, M., Simancik, F.: The incredible ELK - from polynomial procedures to efficient reasoning with el ontologies. *J. Autom. Reason.* **53**(1), 1–61 (2014)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR (Poster) (2015)*
15. Kulmanov, M., Liu-Wei, W., Yan, Y., Hoehndorf, R.: EL embeddings: Geometric construction of models for the description logic EL++. In: *IJCAI*, pp. 6103–6109. ijcai.org (2019)
16. Kulmanov, M., Smaili, F.Z., Gao, X., Hoehndorf, R.: Semantic similarity and machine learning with ontologies. *Briefings Bioinform.* **22**(4) (2021)

⁴ <https://github.com/Box-EL/BoxEL>.

⁵ <https://arxiv.org/abs/2201.09919>.

17. Li, X., Vilnis, L., Zhang, D., Boratko, M., McCallum, A.: Smoothing the geometry of probabilistic box embeddings. In: ICLR. OpenReview.net (2019)
18. Lin, Y., Liu, Z., Sun, M., Liu, Y., Zhu, X.: Learning entity and relation embeddings for knowledge graph completion. In: AAAI, pp. 2181–2187. AAAI Press (2015)
19. Liu, H., Wu, Y., Yang, Y.: Analogical inference for multi-relational embeddings. In: ICML. Proceedings of Machine Learning Research, vol. 70, pp. 2168–2178. PMLR (2017)
20. Mondal, S., Bhatia, S., Mutharaju, R.: Emel++: Embeddings for EL++ description logic. In: AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering. CEUR Workshop Proceedings, vol. 2846. CEUR-WS.org (2021)
21. Mungall, C.J., Torniai, C., Gkoutos, G.V., Lewis, S.E., Haendel, M.A.: Uberon, an integrative multi-species anatomy ontology. *Genome Biol.* **13**(1), 1–20 (2012)
22. Nickel, M., Tresp, V., Kriegel, H.: A three-way model for collective learning on multi-relational data. In: ICML, pp. 809–816. Omnipress (2011)
23. Özçep, Ö.L., Leemhuis, M., Wolter, D.: Cone semantics for logics with negation. In: IJCAI, pp. 1820–1826. ijcai.org (2020)
24. Patel, D., Dasgupta, S.S., Boratko, M., Li, X., Vilnis, L., McCallum, A.: Representing joint hierarchies with box embeddings. In: Automated Knowledge Base Construction (2020). <https://openreview.net/forum?id=J246NSqR.1>
25. Rector, A.L., Rogers, J.E., Pole, P.: The galen high level ontology. In: Medical Informatics Europe'96, pp. 174–178. IOS Press (1996)
26. Ren, H., Hu, W., Leskovec, J.: Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In: ICLR, OpenReview.net (2020)
27. Ren, H., Leskovec, J.: Beta embeddings for multi-hop logical reasoning in knowledge graphs. In: Neurips (2020)
28. Smaili, F.Z., Gao, X., Hoehndorf, R.: Onto2Vec: joint vector-based representation of biological entities and their ontology-based annotations. *Bioinformatics* **34**(13), i52–i60 (2018)
29. Smaili, F.Z., Gao, X., Hoehndorf, R.: Opa2vec: combining formal and informal content of biomedical ontologies to improve similarity-based prediction. *Bioinformatics* **35**(12), 2133–2140 (2019)
30. Steigmiller, A., Liebig, T., Glimm, B.: Konclude: system description. *J. Web Semant.* **27–28**, 78–85 (2014)
31. Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., Bouchard, G.: Complex embeddings for simple link prediction. In: ICML. JMLR Workshop and Conference Proceedings, vol. 48, pp. 2071–2080. JMLR.org (2016)
32. Vilnis, L., Li, X., Murty, S., McCallum, A.: Probabilistic embedding of knowledge graphs with box lattice measures. In: ACL (1), pp. 263–272. Association for Computational Linguistics (2018)
33. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: a survey of approaches and applications. *IEEE Trans. Knowl. Data Eng.* **29**(12), 2724–2743 (2017)
34. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. In: AAAI, pp. 1112–1119. AAAI Press (2014)
35. Yang, B., Yih, W., He, X., Gao, J., Deng, L.: Embedding entities and relations for learning and inference in knowledge bases. In: ICLR (Poster) (2015)