



Dual Box Embeddings for the Description Logic \mathcal{EL}^{++}

Mathias Jackermeier

University of Oxford
Oxford, UK

mathias.jackermeier@cs.ox.ac.uk

Jiaoyan Chen

The University of Manchester
Manchester, UK

jiaoyan.chen@manchester.ac.uk

Ian Horrocks

University of Oxford
Oxford, UK

ian.horrocks@cs.ox.ac.uk

ABSTRACT

OWL ontologies, whose formal semantics are rooted in Description Logic (DL), have been widely used for knowledge representation. Similar to Knowledge Graphs (KGs), ontologies are often incomplete, and maintaining and constructing them has proved challenging. While classical deductive reasoning algorithms use the precise formal semantics of an ontology to predict missing facts, recent years have witnessed growing interest in *inductive* reasoning techniques that can derive *probable* facts from an ontology. Similar to KGs, a promising approach is to learn ontology embeddings in a latent vector space, while additionally ensuring they adhere to the semantics of the underlying DL. While a variety of approaches have been proposed, current ontology embedding methods suffer from several shortcomings, especially that they all fail to faithfully model one-to-many, many-to-one, and many-to-many relations and role inclusion axioms. To address this problem and improve ontology completion performance, we propose a novel ontology embedding method named Box²EL for the DL \mathcal{EL}^{++} , which represents both concepts and roles as boxes (i.e., axis-aligned hyperrectangles), and models inter-concept relationships using a bumping mechanism. We theoretically prove the soundness of Box²EL and conduct an extensive experimental evaluation, achieving state-of-the-art results across a variety of datasets on the tasks of subsumption prediction, role assertion prediction, and approximating deductive reasoning.¹

CCS CONCEPTS

• **Information systems** → *Web Ontology Language (OWL)*; • **Computing methodologies** → **Description logics**; **Ontology engineering**; **Learning latent representations**; *Statistical relational learning*.

KEYWORDS

Ontology Embedding, Ontology Completion, Description Logic, Web Ontology Language, Link Prediction

ACM Reference Format:

Mathias Jackermeier, Jiaoyan Chen, and Ian Horrocks. 2024. Dual Box Embeddings for the Description Logic \mathcal{EL}^{++} . In *Proceedings of the ACM Web Conference 2024 (WWW '24)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3589334.3645648>

¹Code and data are available at <https://github.com/KRR-Oxford/BoxSquaredEL>. Appendices are included in the arXiv version of this paper: <https://arxiv.org/abs/2301.11118>



This work is licensed under a Creative Commons Attribution International 4.0 License.

WWW '24, May 13–17, 2024, Singapore, Singapore
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0171-9/24/05
<https://doi.org/10.1145/3589334.3645648>

1 INTRODUCTION

Ontologies are a widely used formalism to represent general knowledge about a domain [37]. They are usually specified in the Web Ontology Language (OWL) [13] standard developed by W3C², and have been widely adopted in many domains such as the Semantic Web [16], healthcare [34], bioinformatics [17], and geoinformatics [46]. OWL allows for the expression of a variety of statements, ranging from simple relational facts to specifying concept hierarchies and complex logical relationships, and is underpinned by Description Logic (DL) [5] to define its formal semantics.

Although many real-world OWL ontologies have been developed and used with great success, such as the Gene Ontology GO [2] and the food ontology FoodOn [10], both maintaining these existing OWL ontologies and creating new ontologies has proved challenging and relies mostly on manual labor carried out by experts. Common ontology curation tasks include completing missing subsumptions between concepts (or membership relations between individuals and concepts) and identifying missing logical restrictions between concepts. Symbolic logical reasoning algorithms such as HermiT [12] and ELK [19] help address this problem by deductively inferring implicit knowledge from the precise semantics of an ontology, but this classical reasoning is often too rigid for real-world OWL ontologies — especially in the presence of incomplete or noisy data — and cannot derive knowledge that is only *probable* from the given ontology.

At the same time, there has been growing interest in representation learning-based methods for completing Knowledge Graphs (KGs) [18], i.e. relational facts in the form of RDF³ triples $\langle \text{Subject}, \text{Predicate}, \text{Object} \rangle$. Most of these approaches first learn structure-preserving *embeddings* of the entities and relations (predicates) of a KG in a latent vector space and then use them to score the likelihood of novel facts [40]. For example, the classic method TransE [7] maps entities and relations to vectors such that translating the subject embedding by the relation embedding approximately yields the object embedding.

Similar embedding-based techniques as for KGs have been developed for inductive reasoning in ontologies, which promises to complement classical deductive reasoning for ontology curation tasks. Some approaches such as OPA2Vec [35] and OWL2Vec* [9] rely on exploiting textual meta information (e.g., concept labels and comments) to model similarities between entities, but do not retain the semantics defined by the underlying DL. Other approaches aim to directly embed the logical information of an OWL ontology in the latent space [22, 25, 31, 44], mostly targeting the OWL 2 EL profile [21], whose semantics are defined according to the DL \mathcal{EL}^{++} [3]. Prominent examples include ELEM [22] and its extension EmEL⁺⁺ [25], which model concepts as high-dimensional balls, but

²<https://www.w3.org/OWL/>

³Resource Description Framework. <https://www.w3.org/RDF/>.

fail to faithfully capture concept conjunction, since the intersection of two balls is no longer a ball. This led to the development of the state-of-the-art methods BoxEL [44] and ELBE [31], which instead represent concepts as boxes (i.e., axis-aligned hyperrectangles). However, all of these approaches still rely on TransE [7] to model roles (i.e., binary relations) as simple translations, which is unable to capture one-to-many, many-to-one, or many-to-many relationships [23, 42], and is limited in its ability to faithfully represent inclusion relationships between roles. Furthermore, these current works focus on the basic task of predicting subsumptions between named concepts, without considering complex concepts defined with logical operators or complex logical relationships in evaluation.

In this paper, we propose Box²EL, a novel OWL ontology embedding method targeting the semantics of \mathcal{EL}^{++} , which has been widely adopted in many real-life large-scale ontologies [17, 36].⁴ To address the aforementioned limitations of existing approaches, we instead draw inspiration from BoxE [1] and represent both relations and concepts as boxes, while modeling interactions between concepts via a bumping mechanism. We not only demonstrate how Box²EL overcomes the shortcomings of previous methods, but also prove that it is *sound*, i.e., faithfully captures the semantics of the underlying DL, which shows its theoretical correctness and supports interpretable inference for ontology completion. We evaluate our method in the two different inductive reasoning settings of concept subsumption prediction – involving both named and complex concepts – and role assertion (link) prediction, and on approximating deductive reasoning. Our results demonstrate that the theoretical advantages of our approach manifest themselves in practice and lead to state-of-the-art performance across a variety of datasets.

2 BACKGROUND AND RELATED WORK

2.1 Description Logic Ontologies

A DL ontology \mathcal{O} describes some domain of interest in terms of individuals, concepts and roles, where individuals correspond to objects in the domain, concepts represent sets of objects, and roles are binary relations between objects. We limit our discussion to the DL \mathcal{EL}^{++} [3], which underpins the OWL 2 EL profile [21]. It is widely adopted since it contains many useful and important knowledge representation features, while allowing for reasoning in polynomial time. Given sets \mathcal{N}_I , \mathcal{N}_C and \mathcal{N}_R of, respectively, individual, concept, and role names, \mathcal{EL}^{++} concepts are recursively defined as

$$\top \mid \perp \mid A \mid C \sqcap D \mid \exists r.C \mid \{a\}$$

where \top is the top concept, \perp is the bottom concept, $A \in \mathcal{N}_C$ is an atomic (or *named*) concept, $r \in \mathcal{N}_R$ is an atomic role, $a \in \mathcal{N}_I$ is an individual, and C and D are themselves (possibly complex) \mathcal{EL}^{++} concepts. We say a concept is *complex* when it is constructed with a logical operator such as \sqcap or \exists . An \mathcal{EL}^{++} ontology \mathcal{O} consists of a TBox \mathcal{T} and an ABox \mathcal{A} . The TBox consists of logical background knowledge in the form of concept subsumption axioms $C \sqsubseteq D$ and role inclusion axioms $r_1 \circ \dots \circ r_k \sqsubseteq r$, while the ABox contains concrete data in the form of concept and role assertion axioms $C(a)$

⁴The logical constructors provided by \mathcal{EL}^{++} are very common. While some OWL ontologies use more complicated features not supported by \mathcal{EL}^{++} , Box²EL can still be used in that case to model the subset of axioms that fall into \mathcal{EL}^{++} .

and $r(a, b)$. Note that a relational fact from a KG in the form of an RDF triple (a, r, b) is equivalent to a role assertion axiom in the form of $r(a, b)$, and thus ontologies can be seen as extending KGs with more complex conceptual and logical information.

Example 1. The following ontology models a simple family domain:

$$\begin{aligned} \mathcal{T} = \{ & \text{Father} \sqsubseteq \text{Parent} \sqcap \text{Male}, \text{Mother} \sqsubseteq \text{Parent} \sqcap \text{Female}, \\ & \text{Child} \sqsubseteq \exists \text{hasParent.Father}, \text{Child} \sqsubseteq \exists \text{hasParent.Mother}, \\ & \text{hasParent} \sqsubseteq \text{relatedTo} \} \\ \mathcal{A} = \{ & \text{Father}(\text{Alex}), \text{Child}(\text{Bob}), \text{hasParent}(\text{Bob}, \text{Alex}) \} \end{aligned}$$

The TBox specifies that a father is a male parent, a mother is a female parent, every child has a father and a mother, and having a parent implies being related to that parent; the ABox states that Alex is a father, Bob is a child, and Alex is a parent of Bob.

Similarly to first order logic, the semantics of \mathcal{EL}^{++} are defined in terms of *interpretations* that map individuals to elements, concept names to subsets, and role names to binary relations over some set called the *interpretation domain*. An interpretation \mathcal{I} that satisfies the semantics of every axiom in \mathcal{O} is called a *model* of \mathcal{O} , denoted as $\mathcal{I} \models \mathcal{O}$. See Appendix A for a formal discussion of the semantics.

Remark. \mathcal{EL}^{++} also allows for so-called *concrete domains* (a.k.a. datatypes and values), which we do not consider in this paper. Technically, we work on the $\mathcal{ELHO}(\circ)^{\perp}$ subset of \mathcal{EL}^{++} .

2.2 Subsumption Inference

A central problem in DLs is to infer concept subsumptions from an ontology \mathcal{O} . In general, these subsumptions can involve both named and complex concepts. Classical reasoning algorithms leverage the logical information in \mathcal{O} to derive subsumptions that logically follow from the semantics; for example, we can infer from the ontology in Example 1 that $\text{Child} \sqsubseteq \exists \text{relatedTo.Father}$. In contrast, *inductive reasoning* (also called *prediction*) aims to infer *probable* subsumptions from \mathcal{O} . Note that when we limit ourselves to predicting subsumptions of the form $\{a\} \sqsubseteq \exists r.\{b\}$, which are equivalent to role assertion axioms $r(a, b)$, this is identical to the problem of *link prediction* in KGs (i.e., KG completion).

The majority of the existing prediction methods focus only on subsumptions between named concepts. Some approaches embed the formal semantics defined by the DL, while others focus on utilizing textual information such as concept labels and comments. For the former, please see Section 2.4. For the latter, OPA2Vec [35] and OWL2Vec* [9] use a Word2Vec model trained with local graph structure augmented corpora to embed the text for predicting subsumptions, while the recent method BERTSubs [8] fine-tunes a BERT model together with an attached classifier for predicting subsumptions involving both named concepts and complex concepts. Although these works consider a small part of the formal semantics such as the concept hierarchy as the context of a concept for augmenting prediction, they do not model the (complete) logical relationships of the ontology in the vector space. They are complementary to semantic ontology embedding methods including Box²EL, but jointly embedding DL semantics and textual information is out of the scope of this paper.

2.3 Knowledge Graph Embeddings

KG embedding models such as TransE [7], DistMult [45], ComplEx [39] and BoxE [1] aim to solve the problem of completing KGs composed of purely relational facts, and can be thought of as modeling only the role assertion part of the ABox of an OWL ontology [40]. In particular, BoxE [1] also represents relations as boxes and adopts bump vectors, but it models relational facts alone, whereas we aim at much more complex DL ontologies with logical relationships involving concepts and roles.

Some KG embedding methods take background knowledge into account and are therefore related to ontology embedding techniques. However, these methods still focus only on modeling relational facts in a KG, using the background knowledge as constraints, and most of them only support logical rules concerning relations [14, 27, 28, 33, 41] or schemas in simple languages like RDF Schema [15, 43]. In contrast, ontology embedding methods including Box²EL focus on OWL ontologies, which contain a large quantity of conceptual knowledge in the form of subsumptions and logical relationships. Furthermore, these methods can only be applied in the setting of link prediction, whereas ontology embeddings also allow predicting novel conceptual information or logical background knowledge itself.

2.4 Semantic Ontology Embeddings

Several ontology embedding methods for DL semantics have been proposed by learning geometric models. ELEm [22] is among the first to embed \mathcal{EL}^{++} , and EmEL⁺⁺[25] extends ELEm by considering role inclusion axioms. However, both methods represent concepts as high-dimensional balls, which have the disadvantage of not being closed under intersection. Our concept representation based on boxes has previously been used in the two recent methods BoxEL [44] and ELBE [31]. Mondal et al. [25] is the only other technique we are aware of that also models role inclusion axioms; the other methods consider only a smaller subset of \mathcal{EL}^{++} . All previous methods simply model roles (binary relations) by a single vector-based translation as in TransE [7], which fails to faithfully capture one-to-many, many-to-one, and many-to-many relations [23, 42]. We use box-based modeling in combination with bump vectors to address the above problem, achieving better performance in ontology completion.

Going beyond \mathcal{EL}^{++} , Özçep et al. [29] introduce a cone-based model for the more expressive DL \mathcal{ALC} ; however, their contribution is mainly theoretical since they provide neither an implementation nor an evaluation. Embed2Reason [11] is another embedding approach for \mathcal{ALC} ontologies based on quantum logic [6]. In contrast to our work, its focus is on ABox instead of subsumption reasoning. Finally, it is worth mentioning that in comparison with all these DL embedding works, we conduct a more thorough evaluation by considering predicting subsumptions between not only named concepts, but also named concepts and complex concepts involving logical operators.

3 METHOD

In order to perform inductive reasoning over an \mathcal{EL}^{++} ontology \mathcal{O} with signature $\Sigma = (N_C, N_R, N_I)$, we follow the general framework of Kulmanov et al. [22] and learn embeddings that correspond to

geometric models of \mathcal{O} ; that is, (logical) models with interpretation domain $\Delta = \mathbb{R}^n$. We now specify how Box²EL maps concepts, individuals, and roles to the embedding space \mathbb{R}^n and describe the loss functions that encode the axioms of \mathcal{O} .

3.1 Geometric Representation

Concepts and individuals. We follow recent work [31, 44] and represent concepts as n -dimensional *boxes*, i.e., *axis-aligned hyperrectangles*. This representation has several advantages over the alternative based on n -dimensional *balls*, such as closure under concept intersection, and has been shown to work well in practice. Formally, we associate with every concept $C \in N_C$ two vectors $\mathbf{l}_C \in \mathbb{R}^n$ and $\mathbf{u}_C \in \mathbb{R}^n$ such that $\mathbf{l}_C \leq \mathbf{u}_C$, where \leq is applied element-wise. These vectors form the *lower* and *upper* corner of the box of C , i.e., $\text{Box}(C) = \{ \mathbf{x} \in \mathbb{R}^n \mid \mathbf{l}_C \leq \mathbf{x} \leq \mathbf{u}_C \}$. The *center* of $\text{Box}(C)$ is given by $(\mathbf{l}_C + \mathbf{u}_C)/2$, and its *offset* is $(\mathbf{u}_C - \mathbf{l}_C)/2$.

We represent individuals $a \in N_I$ as *points* $\mathbf{e}_a \in \mathbb{R}^n$ in the embedding space. For notational convenience, nominals (e.g., $\{a\}$) are then formally mapped to boxes with volume 0, i.e., $\mathbf{l}_{\{a\}} = \mathbf{u}_{\{a\}} = \mathbf{e}_a$.

Roles. While most existing \mathcal{EL}^{++} embedding models represent roles (binary relations) via simple translations in the form of TransE [7], we instead follow the idea of the BoxE KG embedding model [1]. That is, we associate every role $r \in N_R$ with a *head box* $\text{Head}(r)$ and a *tail box* $\text{Tail}(r)$. Intuitively, every point in the head box is related via r to every point in the tail box. This representation is made more expressive by introducing translational bump vectors $\text{Bump}(C) \in \mathbb{R}^n$ for every atomic concept C , which model interactions between related concepts by dynamically translating their embeddings. An axiom of the form $C \sqsubseteq \exists r.D$ is then considered to hold if

$$\begin{aligned} \text{Box}(C) \oplus \text{Bump}(D) &\subseteq \text{Head}(r) \quad \text{and} \\ \text{Box}(D) \oplus \text{Bump}(C) &\subseteq \text{Tail}(r), \end{aligned} \tag{1}$$

where \oplus denotes the operation of translating a box along a bump vector, i.e., $\text{Box}(C) \oplus \text{Bump}(D) = \{ \mathbf{x} + \text{Bump}(D) \mid \mathbf{x} \in \text{Box}(C) \}$.

Example 2. Figure 1 illustrates 2-dimensional Box²EL embeddings that form a logical model of the TBox in Example 1, since, e.g., $\text{Box}(\text{Father}) \subseteq \text{Box}(\text{Parent}) \cap \text{Box}(\text{Male})$ and Equation (1) holds for all relevant axioms. We will formalise the logical geometric model associated with a set of Box²EL embeddings in Theorem 1. Note also that the embeddings furthermore imply the subsumption $\text{Child} \sqsubseteq \exists \text{relatedTo.Father}$ (again by Equation (1)), illustrating their utility for approximate reasoning.

Expressiveness. The previous example illustrates the expressive power of Box²EL. Representing roles as head and tail boxes allows modeling one-to-many relationships such as *hasParent* faithfully, in contrast to previous approaches that employ a TransE-based role representation. Moreover, role inclusion axioms such as *hasParent* \sqsubseteq *relatedTo* can be represented naturally via inclusion constraints on the relevant head and tail boxes. In contrast, previous methods either do not consider role inclusion axioms at all or only rudimentarily approximate them by forcing the embeddings of the involved roles to be similar [25].

Model complexity. Box²EL requires $2n|N_C| + n|N_I|$ parameters to store the lower and upper corners of the box embeddings of concepts and points associated with individuals. In order to represent

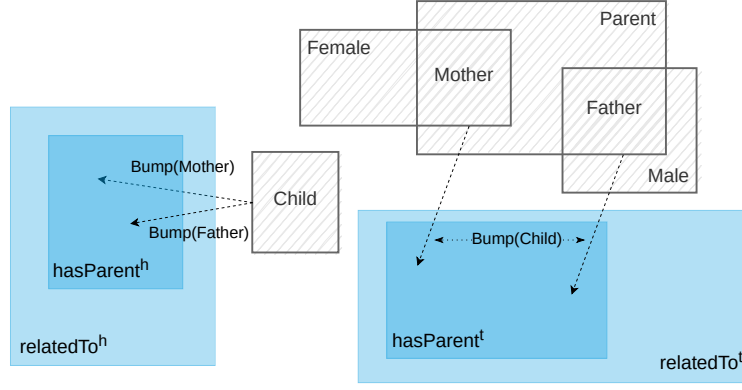


Figure 1: An illustration of Box^2EL embeddings. Striped boxes represent concept embeddings, whereas role embeddings are shaded blue and labelled as r^h or r^t for the head or tail box of r , respectively. Bump vectors are drawn as arrows and labelled with the corresponding concept. The illustrated embeddings form a logical model of the TBox in Example 1.

the head and tail boxes for every relation and a bump vector per concept and individual, the model needs $4n|\mathcal{N}_R| + n(|\mathcal{N}_C| + |\mathcal{N}_I|)$ additional parameters. Therefore, the total space complexity of Box^2EL is in $O(n(3|\mathcal{N}_C| + 2|\mathcal{N}_I| + 4|\mathcal{N}_R|))$.

3.2 Training Procedure

In order to learn embeddings for \mathcal{O} , we first convert its ABox axioms into equivalent TBox axioms using the transformation rules

$$\begin{aligned} C(a) &\rightsquigarrow \{a\} \sqsubseteq C \\ r(a, b) &\rightsquigarrow \{a\} \sqsubseteq \exists r. \{b\} \end{aligned}$$

and then normalize the axioms using the standard procedure in [3, 4]. This results in a normalized ontology in which all axioms are in one of the normal forms described below. Crucially, this procedure is semantics-preserving: every model of the normalized ontology is also a model of the original ontology [3]. For more details on normalization, see Appendix B. We introduce separate loss functions for axioms in each normal form, which intuitively ensure that the learned embeddings adhere to the semantics of \mathcal{O} . Finally, we minimize the sum of all loss terms $\mathcal{L}(\mathcal{O})$ via mini-batch gradient descent.

Our loss functions are based on the distance-based loss formulation in [22, 31] and aim to minimize the element-wise distance between the embeddings of related concepts. Given two arbitrary boxes A and B , this element-wise distance is computed as

$$d(A, B) = |c(A) - c(B)| - o(A) - o(B),$$

where $c(\cdot)$ and $o(\cdot)$ denote the center and offset of a box, respectively. Note that for nominals $\{a\}$ we have that $c(\text{Box}(\{a\})) = e_a$ and $o(\text{Box}(\{a\})) = \mathbf{0}$.

Generic inclusion loss. Before defining loss functions for the different \mathcal{EL}^{++} normal forms, we first introduce a generic inclusion loss $\mathcal{L}_{\subseteq}(A, B)$. It encourages box A to be contained in box B and is defined as

$$\mathcal{L}_{\subseteq}(A, B) = \begin{cases} \|\max\{\mathbf{0}, d(A, B) + 2o(A) - \gamma\}\| & \text{if } B \neq \emptyset \\ \max\{0, o(A)_1 + 1\} & \text{otherwise,} \end{cases}$$

where $\gamma \in \mathbb{R}$ is a margin hyperparameter. If $\mathcal{L}_{\subseteq}(A, B) = 0$, either A lies within γ -distance of B in each dimension, or both A and B are empty.

We next introduce each normal form and the corresponding loss function. Note that all concepts in the normal forms below are atomic concepts or nominals (and not complex).

First normal form (NF1). For an NF1 axiom of the form $C \sqsubseteq D$, the learned embeddings need to satisfy $\text{Box}(C) \subseteq \text{Box}(D)$, corresponding to the semantics of concept inclusion. Therefore, we define the loss for NF1 as simply the inclusion loss:

$$\mathcal{L}_1(C, D) = \mathcal{L}_{\subseteq}(\text{Box}(C), \text{Box}(D)).$$

Second normal form (NF2). For an NF2 axiom of the form $C \sqcap D \sqsubseteq E$, we similarly require that the intersection of the boxes of C and D is within the box of E . The intersection of $\text{Box}(C)$ and $\text{Box}(D)$ is itself a box with lower corner $\max\{l_C, l_D\}$ and upper corner $\min\{u_C, u_D\}$, where \max and \min are applied element-wise. We thus have

$$\mathcal{L}_2(C, D, E) = \mathcal{L}_{\subseteq}(\text{Box}(C) \cap \text{Box}(D), \text{Box}(E)).$$

However, this formulation is problematic since it can be easily minimized to 0 by setting $\text{Box}(C)$ and $\text{Box}(D)$ to be disjoint. While disjoint embeddings for C and D would technically not violate the semantics, usually an axiom of the form $C \sqcap D \sqsubseteq \perp$ would have been used directly if it had been the intention that C and D should be disjoint. Therefore, we add the additional term

$$\|\max\{\mathbf{0}, \max\{l_C, l_D\} - \min\{u_C, u_D\}\}\|$$

to the loss, which intuitively encourages $\text{Box}(C) \cap \text{Box}(D)$ to be non-empty by making all elements of its offset vector positive.

Third normal form (NF3). For an NF3 axiom of the form $C \sqsubseteq \exists r.D$, the embeddings should satisfy $\text{Box}(C) + \text{Bump}(D) \subseteq \text{Head}(r)$ and $\text{Box}(D) + \text{Bump}(C) \subseteq \text{Tail}(r)$. This is captured by the following loss function:

$$\begin{aligned} \mathcal{L}_3(C, r, D) = \frac{1}{2} &\left(\mathcal{L}_{\subseteq}(\text{Box}(C) + \text{Bump}(D), \text{Head}(r)) \right. \\ &\left. + \mathcal{L}_{\subseteq}(\text{Box}(D) + \text{Bump}(C), \text{Tail}(r)) \right). \end{aligned}$$

Father \sqsubseteq Male \sqcap Parent	Mother \sqsubseteq Female \sqcap Parent
Male \sqcap Parent \sqsubseteq Father	Female \sqcap Parent \sqsubseteq Mother
Male \sqcap Female $\sqsubseteq \perp$	Parent \sqcap Child $\sqsubseteq \perp$
Child $\sqsubseteq \exists \text{hasParent.Mother}$	Child $\sqsubseteq \exists \text{hasParent.Father}$
Parent $\sqsubseteq \exists \text{hasChild.Child}$	

Figure 2: Proof of concept ontology.

If $\text{Box}(D)$ is empty, we furthermore add the term $\mathcal{L}_{\subseteq}(\text{Box}(C), \emptyset)$ to also make $\text{Box}(C)$ empty.

Fourth normal form (NF4). For an NF4 axiom of the form $\exists r.C \sqsubseteq D$, we need to ensure that all points in the embedding space that are connected to C via role r are contained in $\text{Box}(D)$. It can be seen from our geometric representation that the set of these points is contained in the set $\text{Head}(r) - \text{Bump}(C)$. We therefore define the loss for the fourth normal form as

$$\mathcal{L}_4(r, C, D) = \mathcal{L}_{\subseteq}(\text{Head}(r) - \text{Bump}(C), \text{Box}(D)).$$

Fifth normal form (NF5). Axioms of the fifth normal form $C \sqcap D \sqsubseteq \perp$ state that concepts C and D are disjoint. Our corresponding loss function penalizes embeddings for which the element-wise distance is not greater than 0 (within a margin of γ) and is defined as

$$\mathcal{L}_5(C, D) = \|\max\{\mathbf{0}, -(\mathbf{d}(\text{Box}(C), \text{Box}(D)) + \gamma)\}\|.$$

Role inclusion axioms. After normalization, role inclusion axioms are either of the form $r \sqsubseteq s$ or $r_1 \circ r_2 \sqsubseteq s$. For the first case, we define the loss function

$$\mathcal{L}_6(r, s) = \frac{1}{2} \left(\mathcal{L}_{\subseteq}(\text{Head}(r), \text{Head}(s)) + \mathcal{L}_{\subseteq}(\text{Tail}(r), \text{Tail}(s)) \right),$$

which intuitively makes any embeddings related via role r also related via s . Similarly, in the second case we have the loss

$$\mathcal{L}_7(r_1, r_2, s) = \frac{1}{2} \left(\mathcal{L}_{\subseteq}(\text{Head}(r_1), \text{Head}(s)) + \mathcal{L}_{\subseteq}(\text{Tail}(r_2), \text{Tail}(s)) \right).$$

Regularization. In order to prevent our expressive role representation from overfitting, we furthermore L^2 -regularize the bump vectors by adding the regularization term

$$\lambda \sum_{C \in \mathcal{N}_C \cup \mathcal{N}_I} \|\text{Bump}(C)\|,$$

where λ is a hyperparameter.

Negative sampling. In addition to the above loss functions, we also employ *negative sampling* to further improve the quality of the learned embeddings. We follow previous work [22, 31] and generate negative samples for axioms of the form $C \sqsubseteq \exists r.D$ by replacing either C or D with a randomly selected different concept, similar to negative sampling in KGs [7]. For every NF3 axiom we generate a new set of $\omega \geq 1$ negative samples $C \not\sqsubseteq \exists r.D$ in every epoch, for which we optimize the loss

$$\begin{aligned} \mathcal{L}_{\not\subseteq}(C, r, D) &= (\delta - \mu(\text{Box}(C) + \text{Bump}(D), \text{Head}(r)))^2 \\ &\quad + (\delta - \mu(\text{Box}(D) + \text{Bump}(C), \text{Tail}(r)))^2, \end{aligned}$$

where $\mu(A, B) = \|\max\{\mathbf{0}, \mathbf{d}(A, B) + \gamma\}\|$ is the minimal distance between any two points in A and B and $\delta > 0$ is a hyperparameter that controls how unlikely the negative samples are made by the model.

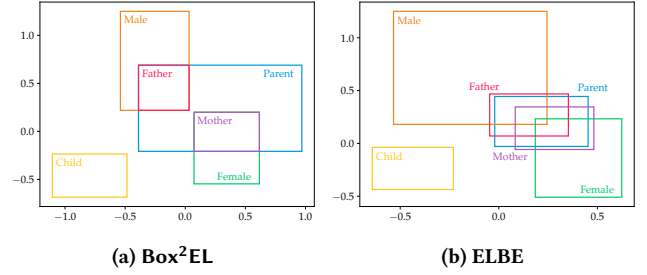


Figure 3: Visualization of the embeddings learned by Box²EL (left) and ELBE (right) for the proof-of-concept ontology. While Box²EL can accurately represent the axioms in the ontology, the limitations of TransE as a model for roles prevent ELBE from learning correct embeddings.

As in the case of KGs, the above procedure may occasionally generate false negative samples, i.e., negative axioms that actually do follow from \mathcal{O} . However, on average it will predominantly produce true negatives, which we find to empirically improve the learned embeddings.

3.3 Soundness

We now show that the embeddings learned by Box²EL indeed correspond to a logical geometric model of the given ontology \mathcal{O} .

THEOREM 1 (SOUNDNESS). *Let $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ be an \mathcal{EL}^{++} ontology. If $\gamma \leq 0$ and there exist Box²EL embeddings in \mathbb{R}^n such that $\mathcal{L}(\mathcal{O}) = 0$, then these embeddings are a model of \mathcal{O} .*

PROOF SKETCH. The proof relies on the correctness of our loss functions, i.e., if $\mathcal{L}_{\subseteq}(A, B) = 0$ then $A \subseteq B$. We first construct a geometric interpretation \mathcal{I} of \mathcal{O} from the embeddings by interpreting every individual as its associated point, every concept as the set of points in its associated box, and every role as the Cartesian product of its head box and its tail box. Since the loss for every axiom is 0, the embeddings satisfy the corresponding semantics. We thus have that \mathcal{I} is a model of \mathcal{O} . \square

A formal proof of Theorem 1 is given in Appendix C. While our optimization procedure might not achieve a loss of 0 in practice, the importance of Theorem 1 is that it demonstrates that the learned embeddings converge to a semantically meaningful representation of the ontology in which all of its axioms are satisfied, i.e., a model of \mathcal{O} . The embeddings therefore indeed encode the axioms in \mathcal{O} and are thus useful to perform inductive or approximate reasoning.

4 EVALUATION

We first validate our model and demonstrate its expressiveness on a proof of concept ontology. We then evaluate Box²EL on three tasks: general subsumption prediction, link (role assertion) prediction, and approximating deductive reasoning. Furthermore, we conduct a variety of ablation studies whose results are shown in Appendix E. Our implementation is based on PyTorch [30], and we use the jcel reasoner [24] to transform ontologies into normal form axioms.

Table 1: Overall subsumption prediction results combined across normal forms.

	Model	H@1	H@10	H@100	Med	MRR	MR	AUC
GALEN	ELEm	0.01	0.12	0.29	1662	0.05	5153	0.78
	EmEL ⁺⁺	0.01	0.11	0.24	2295	0.05	5486	0.76
	BoxEL	0.00	0.03	0.16	4785	0.01	7163	0.69
	ELBE	0.02	0.14	0.27	1865	0.06	5303	0.77
	Box ² EL	0.05	0.20	0.35	669	0.10	4375	0.81
GO	ELEm	0.03	0.24	0.43	272	0.09	6204	0.86
	EmEL ⁺⁺	0.03	0.23	0.38	597	0.09	6710	0.85
	BoxEL	0.01	0.06	0.08	8443	0.03	14905	0.68
	ELBE	0.01	0.10	0.22	1838	0.04	6986	0.85
	Box ² EL	0.04	0.23	0.59	48	0.10	3248	0.93
Anatomy	ELEm	0.10	0.40	0.64	22	0.19	6464	0.94
	EmEL ⁺⁺	0.11	0.36	0.57	36	0.19	8472	0.92
	BoxEL	0.03	0.12	0.28	1151	0.06	10916	0.90
	ELBE	0.04	0.36	0.63	29	0.15	5400	0.95
	Box ² EL	0.16	0.47	0.70	13	0.26	2675	0.97

4.1 Proof of Concept: Family Ontology

We visualize the embeddings learned by Box²EL for a proof of concept family ontology given in Figure 2. To this end, we train Box²EL with an embedding dimensionality of $n = 2$, a margin of $\gamma = 0$, regularization strength $\lambda = 1$, and no negative sampling. We also train embeddings for ELBE [31], a comparable \mathcal{EL}^{++} ontology embedding model that similarly uses boxes to represent concepts, but interprets roles as translations. In order to ensure the volume of the learned embeddings is large enough for plotting, we add the following visualization loss term to the objective function of both models:

$$\mathcal{L}_V = \frac{1}{n|\mathcal{N}_C|} \sum_{C \in \mathcal{N}_C} \sum_{i=1}^n \max\{0, 0.2 - \sigma(\text{Box}(C))_i\}.$$

The learned concept embeddings of both models are depicted in Figure 3. We see that Box²EL is able to successfully learn embeddings that align with the axioms in the ontology. In particular, the embeddings fulfill all disjointness axioms and correctly represent the relationships between the concepts Father, Male, Mother, Female, and Parent.

In contrast, we find that the embeddings learned by ELBE violate several of the axioms in the ontology. This is due to the inability of the underlying TransE model to correctly represent one-to-many relationships: because the ontology contains the axioms $\text{Child} \sqsubseteq \exists \text{hasParent.Mother}$ as well as $\text{Child} \sqsubseteq \exists \text{hasParent.Father}$, the model is forced to let the embeddings of Mother and Father overlap.

4.2 General Subsumption Prediction

We next evaluate Box²EL on general subsumption prediction for inductive reasoning. In contrast to previous work [25, 44], we not only consider subsumptions between atomic (named) concepts, but also the more challenging task of predicting subsumptions between atomic concepts and complex concepts.

Benchmark. We introduce a modification of the benchmark based on three biomedical ontologies GALEN [32], Gene Ontology (GO) [2]

Table 2: Detailed subsumption prediction results on the GALEN ontology.

	Model	H@1	H@10	H@100	Med	MRR	MR	AUC
NF1	ELEm	0.01	0.16	0.40	430	0.06	3568	0.85
	EmEL ⁺⁺	0.02	0.16	0.37	632	0.06	3765	0.84
	BoxEL	0.00	0.00	0.05	3715	0.00	5727	0.75
	ELBE	0.03	0.24	0.47	138	0.10	2444	0.89
	Box ² EL	0.03	0.30	0.51	91	0.12	2632	0.89
NF2	ELEm	0.01	0.07	0.17	5106	0.03	7432	0.68
	EmEL ⁺⁺	0.01	0.07	0.15	5750	0.03	7767	0.66
	BoxEL	0.00	0.00	0.00	11358	0.00	11605	0.50
	ELBE	0.03	0.06	0.11	6476	0.04	8068	0.65
	Box ² EL	0.06	0.15	0.28	2149	0.09	6265	0.73
NF3	ELEm	0.02	0.14	0.28	1479	0.05	4831	0.79
	EmEL ⁺⁺	0.02	0.11	0.22	2240	0.05	5348	0.77
	BoxEL	0.00	0.02	0.08	7239	0.01	8615	0.63
	ELBE	0.03	0.14	0.25	2154	0.07	5072	0.78
	Box ² EL	0.08	0.19	0.32	635	0.12	3798	0.84
NF4	ELEm	0.00	0.05	0.18	3855	0.02	6793	0.71
	EmEL ⁺⁺	0.00	0.04	0.12	4458	0.01	7020	0.70
	BoxEL	0.00	0.15	0.69	47	0.04	2667	0.89
	ELBE	0.00	0.03	0.07	7563	0.01	8884	0.62
	Box ² EL	0.00	0.06	0.15	4364	0.02	7266	0.69

and Anatomy (a.k.a. Uberon) [26], that has been considered in previous DL embedding works [25, 44]. For each ontology, our benchmark consists of axioms split into training (80%), validation (10%), and testing (10%) sets for each normal form. This enables us to evaluate ontology embedding models on subsumption prediction between atomic concepts (NF1), atomic concepts and conjunctions (NF2), and atomic concepts and existential restrictions (NF3 and NF4). We report statistics on the sizes of these ontologies, including the numbers of axioms of different forms, in Table 7 in the appendix.

Baselines. We compare Box²EL with the state-of-the-art ontology embedding methods ELEm [22], EmEL⁺⁺ [25], BoxEL [44], and ELBE [31]. We do not consider any traditional KG embedding methods in our experiments, since they have been shown to be considerably outperformed by ontology embedding methods [25, 44] and are not applicable in the setting of complex concepts.

Evaluation protocol. To evaluate the subsumption prediction performance, we follow the literature [25, 44] and report a variety of ranking-based metrics on the testing set. Given a test axiom in some normal form, we generate a set of candidate predictions by replacing the atomic side of the subsumption with all the atomic concepts in \mathcal{N}_C . We then rank all candidate predictions by a score based on the distance between the embeddings of the concepts of the subsumption (for details see Appendix F) and record the rank of the true axiom. We report the standard metrics Hits@ k ($H@k$), where $k \in \{1, 10, 100\}$, the median rank (Med), the mean reciprocal rank (MRR), the mean rank (MR), and the area under the ROC curve (AUC). These metrics are computed for the axioms in each normal form individually, as well as combined across normal forms. See Appendix G for definitions of the metrics.

Experimental protocol. The embeddings are optimized with Adam [20] for a maximum of 10,000 epochs. All hyperparameters are

Table 3: Detailed subsumption prediction results on the GO ontology.

	Model	H@1	H@10	H@100	Med	MRR	MR	AUC
NF1	ELEm	0.01	0.13	0.35	590	0.05	6433	0.86
	EmEL ⁺⁺	0.01	0.12	0.30	1023	0.05	6709	0.85
	BoxEL	0.00	0.01	0.05	5374	0.00	13413	0.71
	ELBE	0.01	0.10	0.24	1156	0.04	5657	0.88
	Box ² EL	0.03	0.17	0.58	58	0.08	2686	0.94
NF2	ELEm	0.12	0.49	0.63	11	0.24	4508	0.90
	EmEL ⁺⁺	0.11	0.44	0.55	23	0.21	5169	0.89
	BoxEL	0.00	0.00	0.00	22882	0.00	23007	0.50
	ELBE	0.01	0.05	0.09	6456	0.02	9421	0.80
	Box ² EL	0.18	0.58	0.75	6	0.31	2104	0.95
NF3	ELEm	0.06	0.40	0.52	54	0.15	6292	0.86
	EmEL ⁺⁺	0.05	0.39	0.48	210	0.15	7788	0.83
	BoxEL	0.00	0.00	0.00	17027	0.00	18947	0.59
	ELBE	0.02	0.15	0.30	959	0.07	7131	0.84
	Box ² EL	0.00	0.18	0.53	79	0.05	5042	0.89
NF4	ELEm	0.01	0.49	0.60	12	0.12	6272	0.86
	EmEL ⁺⁺	0.01	0.49	0.58	12	0.13	6442	0.86
	BoxEL	0.09	0.54	0.54	2215	0.28	9673	0.79
	ELBE	0.00	0.07	0.12	9049	0.02	12868	0.72
	Box ² EL	0.00	0.37	0.64	20	0.08	4989	0.89

described in detail in Appendix H. We evaluate the models on a fraction of the validation set every 100 epochs and choose the embeddings that achieve the best performance for final evaluation on the testing set. The results we report are averages across 5 different random seeds.

Results. The results on all the testing axioms (combined across all normal forms) are reported in Table 1. For detailed results on testing axioms of each normal form, see Tables 2 to 4. We first observe that our model Box²EL consistently outperforms all the baselines on all datasets, often with significant performance gains. For example, the median rank (MR) of Box²EL is around 60% lower than the second best-performing method on GALEN, more than 80% lower on GO, and more than 40% lower on Anatomy. Among the baseline methods, results are similar; interestingly, ELEm generally performs best, in contrast to previous benchmarks.

From the detailed results, we observe that the novel role representation of Box²EL not only generally improves prediction performance for NF3 axioms, which contain roles, but also for NF1 and NF2 axioms. This can be explained by the fact that the different normal forms are used to optimize the *same* embeddings; i.e., if Box²EL can better represent an axiom of the form $C \sqsubseteq \exists r.D$, it will learn better embeddings for C and D , therefore also improving prediction quality for NF1 and NF2 axioms involving C and/or D . There is no clear trend which axioms are the easiest to predict; on GALEN, the models generally perform better on NF1 axioms involving only atomic concepts, whereas on GO and Anatomy they perform similarly well on axioms involving complex concepts.

4.3 Link Prediction

We next evaluate our model on the task of link prediction, i.e., predicting role assertions of the form $r(a, b)$, which is implemented by predicting subsumptions of the form $\{a\} \sqsubseteq \exists r.\{b\}$.

Table 4: Detailed subsumption prediction results on the Anatomy ontology.

	Model	H@1	H@10	H@100	Med	MRR	MR	AUC
NF1	ELEm	0.07	0.30	0.57	43	0.14	9059	0.91
	EmEL ⁺⁺	0.08	0.29	0.53	60	0.14	10414	0.90
	BoxEL	0.01	0.05	0.16	1828	0.03	9597	0.91
	ELBE	0.05	0.24	0.55	68	0.11	5177	0.95
	Box ² EL	0.07	0.34	0.65	27	0.15	2894	0.97
NF2	ELEm	0.03	0.18	0.42	394	0.08	11592	0.89
	EmEL ⁺⁺	0.03	0.18	0.35	1291	0.08	15759	0.85
	BoxEL	0.00	0.00	0.00	17607	0.00	26872	0.75
	ELBE	0.02	0.11	0.26	1394	0.05	4885	0.96
	Box ² EL	0.16	0.41	0.64	26	0.24	1928	0.98
NF3	ELEm	0.12	0.47	0.69	13	0.23	4686	0.96
	EmEL ⁺⁺	0.13	0.42	0.60	23	0.23	7097	0.93
	BoxEL	0.04	0.17	0.36	567	0.08	11095	0.90
	ELBE	0.04	0.44	0.70	16	0.18	5408	0.95
	Box ² EL	0.21	0.56	0.75	7	0.33	2466	0.98
NF4	ELEm	0.00	0.03	0.23	813	0.01	10230	0.91
	EmEL ⁺⁺	0.00	0.02	0.17	1470	0.01	10951	0.90
	BoxEL	0.00	0.00	0.00	38942	0.00	41283	0.61
	ELBE	0.00	0.02	0.06	6261	0.01	15187	0.86
	Box ² EL	0.00	0.05	0.14	3065	0.01	8366	0.92

Datasets. We consider a real-world protein-protein interaction (PPI) prediction task introduced in [22]. They provide two ontologies for human and yeast organisms, constructed by combining the STRING database of PPIs [38] with the Gene Ontology [2]. The proteins and their interactions recorded in STRING constitute the ABox, while GO acts as the TBox, and additional axioms modeling the association of proteins with their biological functions are added to the ontology. The task is to predict missing subsumptions of the form $\{P_1\} \sqsubseteq \exists \text{interacts.}\{P_2\}$, where P_1 and P_2 represent two proteins.

Baselines. We also consider ELEm [22], EmEL⁺⁺ [25], BoxEL [44], and ELBE [31] for the baselines as in general subsumption prediction, and report the relevant best results from their original papers.

Evaluation and experimental protocol. In order to evaluate our method, we use the 80%/10%/10% training, validation, and testing split of the PPI data provided by [22]. We compute the same ranking-based metrics as in subsumption prediction, both in the standard and filtered fashion, in which any true candidate predictions except for the target axiom to predict are first removed from the set of all candidate predictions before computing the ranks. The experimental protocol is the same as in subsumption prediction.

Results. Table 5 lists the results of Box²EL and the baseline methods on the yeast and human PPI prediction datasets. Box²EL outperforms all the baselines, with significant performance gains for both datasets on most of the metrics including filtered hits and mean rank. The AUC and AUC (F) values are all very close to the maximum value 1.0, but Box²EL still improves the state-of-the-art on Yeast, and ties ELBE on Human. All these results demonstrate the effectiveness of Box²EL in role assertion prediction for ontologies with an ABox. The comparatively stronger results than ELBE and BoxEL, which share the same concept representation as our model,

Table 5: PPI prediction results on the yeast and human datasets. Columns annotated with (F) contain filtered metrics, other columns contain raw metrics. The results for BoxEL are from [44]; all other baseline results are from [31].

	Model	H@10	H@10 (F)	H@100	H@100 (F)	MR	MR (F)	AUC	AUC (F)
Yeast	ELEm	0.10	0.23	0.50	0.75	247	187	0.96	0.97
	EmEL ⁺⁺	0.08	0.17	0.48	0.65	336	291	0.94	0.95
	BoxEL	0.09	0.20	0.52	0.73	423	379	0.93	0.94
	ELBE	0.11	0.26	0.57	0.77	201	154	0.96	0.97
	Box ² EL	0.11	0.33	0.64	0.87	168	118	0.97	0.98
Human	ELEm	0.09	0.22	0.43	0.70	658	572	0.96	0.96
	EmEL ⁺⁺	0.04	0.13	0.38	0.56	772	700	0.95	0.95
	BoxEL	0.07	0.10	0.42	0.63	1574	1530	0.93	0.93
	ELBE	0.09	0.22	0.49	0.72	434	362	0.97	0.98
	Box ² EL	0.09	0.28	0.55	0.83	343	269	0.98	0.98

Table 6: Deductive reasoning results on GALEN, GO, and Anatomy.

	Model	H@1	H@10	H@100	Med	MRR	MR	AUC
GALEN	ELEm	0.00	0.04	0.20	1807	0.01	4405	0.81
	EmEL ⁺⁺	0.00	0.04	0.18	2049	0.01	4634	0.81
	BoxEL	0.00	0.00	0.01	6906	0.00	7925	0.67
	ELBE	0.00	0.06	0.16	1785	0.02	3974	0.84
	Box ² EL	0.01	0.09	0.24	1003	0.03	2833	0.88
GO	ELEm	0.00	0.04	0.22	1629	0.02	7377	0.84
	EmEL ⁺⁺	0.00	0.04	0.19	1346	0.01	6557	0.86
	BoxEL	0.00	0.00	0.13	1085	0.00	5359	0.88
	ELBE	0.00	0.06	0.21	935	0.02	3846	0.92
	Box ² EL	0.00	0.08	0.49	107	0.04	1689	0.96
Anatomy	ELEm	0.00	0.07	0.28	901	0.02	7958	0.93
	EmEL ⁺⁺	0.00	0.07	0.26	1576	0.02	10976	0.90
	BoxEL	0.01	0.10	0.24	838	0.04	9156	0.92
	ELBE	0.00	0.08	0.32	336	0.03	2312	0.98
	Box ² EL	0.01	0.09	0.44	152	0.04	1599	0.99

once again highlight the positive impact of our novel approach of representing the semantics of roles with boxes and bump vectors.

4.4 Approximating Deductive Reasoning

We finally investigate how well our model can approximate *deductive reasoning*, i.e., infer subsumptions that are logical consequences of the axioms in the ontology.

Experimental setup. We again consider the GALEN, GO, and Anatomy ontologies. Instead of splitting the axioms into separate training, validation, and testing sets, we now train the models on the entire ontology using all of its asserted axioms. For evaluation, we use an \mathcal{EL}^{++} reasoner to compute the complete set of NF1 axioms (i.e., atomic subsumptions) that are logically implied by, but not explicitly asserted in the given ontology. We split off 10% of these implied NF1 axioms for the validation set and keep the remainder as the testing set. We report the results of Box²EL and the same baseline methods considered in subsumption prediction. The evaluation and experimental protocol is also the same as in subsumption prediction.⁵

⁵Deductive reasoning with \mathcal{EL}^{++} embeddings has been previously considered in [25]. However, we find that there is significant leakage (overlap) between their testing and

Results. Table 6 lists the results of approximating deductive reasoning. The baselines perform similarly, with ELBE achieving slightly stronger results than the others on GO and Anatomy. Box²EL outperforms the baselines on almost all metrics across the three ontologies, with significant performance gains especially for Hits@100, median rank, and mean rank. This indicates that Box²EL is able to preserve more of the logical structure than the other embedding methods.

Deductive vs inductive reasoning. Comparing the results in Tables 1 and 6, we observe that the embedding models generally perform better in subsumption prediction than in approximating deductive reasoning. To see why, note that the learned embeddings are used to make purely statistical predictions about missing axioms. The soundness of our method guarantees that these predictions align with the semantics of the ontology. However, we do not explicitly perform logical inference steps in the embedding space, as would be required to derive logical inferences similar to a deductive reasoning algorithm. We illustrate this difference with a concrete example in Appendix I. While embedding methods can thus be useful to approximate deductive reasoning, the two approaches are best used in conjunction in order to combine formal derivations with inductive and probable knowledge.

5 CONCLUSION

We developed Box²EL, a novel OWL ontology embedding method that adopts box-based representations for both concepts and roles. This representation is able to model complex logical constructs from \mathcal{EL}^{++} and overcomes the limitations of previous approaches in representing roles and role inclusion axioms. We formally proved that our method is sound, i.e., correctly represents the semantics of \mathcal{EL}^{++} , and performed an extensive empirical evaluation, achieving state-of-the-art results in concept subsumption prediction, role assertion prediction, and approximating deductive reasoning.

ACKNOWLEDGMENTS

This work was supported by Samsung Research UK (SRUK) and the EPSRC projects ConCur (EP/V050869/1), OASIS (EP/S032347/1), and UK FIRES (EP/S019111/1).

training sets. We therefore do not adopt their reported results, but instead reproduce the results of the baseline models.

REFERENCES

- [1] Ralph Abboud, İsmail İlkan Ceylan, Thomas Lukasiewicz, and Tommaso Salvatori. 2020. BoxE: A Box Embedding Model for Knowledge Base Completion. In *Advances in Neural Information Processing Systems*, Vol. 33. 9649–9661.
- [2] Michael Ashburner, Catherine A Ball, Judith A Blake, David Botstein, Heather Butler, J Michael Cherry, Allan P Davis, Kara Dolinski, Selina S Dwight, Janan T Eppig, et al. 2000. Gene ontology: tool for the unification of biology. *Nature genetics* 25, 1 (2000), 25–29.
- [3] Franz Baader, Sebastian Brandt, and Carsten Lutz. 2005. Pushing the EL envelope. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, Vol. 5. 364–369.
- [4] Franz Baader, Sebastian Brandt, and Carsten Lutz. 2005. *Pushing the EL Envelope*. LTCS-Report LTCS-05-01. Institute for Theoretical Computer Science, TU Dresden. <http://lat.inf.tu-dresden.de/research/reports.html>
- [5] Franz Baader, Ian Horrocks, and Ulrike Sattler. 2005. Description logics as ontology languages for the semantic web. In *Mechanizing mathematical reasoning*. Springer, 228–248.
- [6] Garrett Birkhoff and John Von Neumann. 1936. The Logic of Quantum Mechanics. *Annals of Mathematics* 37, 4 (1936), 823–843. [jstor:1968621](https://doi.org/10.2307/1968621)
- [7] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems*, Vol. 26.
- [8] Jiaoyan Chen, Yuan He, Yuxia Geng, Ernesto Jiménez-Ruiz, Hang Dong, and Ian Horrocks. 2023. Contextual semantic embeddings for ontology subsumption prediction. *World Wide Web* (2023), 1–23.
- [9] Jiaoyan Chen, Pan Hu, Ernesto Jimenez-Ruiz, Ole Magnus Holter, Denver Antonyrajah, and Ian Horrocks. 2021. OWL2Vec: Embedding of OWL ontologies. *Machine Learning* 110, 7 (2021), 1813–1845.
- [10] Damion M. Dooley, Emma J. Griffiths, Gurinder S. Gosal, Pier L. Buttigieg, Robert Hoehndorf, Matthew C. Lange, Lynn M. Schriml, Fiona S. L. Brinkman, and William W. L. Hsiao. 2018. FoodOn: A Harmonized Food Ontology to Increase Global Food Traceability, Quality Control and Data Integration. *npj Science of Food* 2, 1 (Dec. 2018), 23.
- [11] Dinesh Garg, Shajith Ikkal, Santosh K Srivastava, Harit Vishwakarma, Hima Karanam, and L Venkata Subramaniam. 2019. Quantum embedding of knowledge for reasoning. *Advances in Neural Information Processing Systems* 32 (2019).
- [12] Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. 2014. Hermit: An OWL 2 Reasoner. *Journal of Automated Reasoning* 53, 3 (2014), 245–269.
- [13] Bernardo Cuenca Grau, Ian Horrocks, Boris Motik, Bijan Parsia, Peter Patel-Schneider, and Ulrike Sattler. 2008. OWL 2: The next Step for OWL. *Journal of Web Semantics* 6, 4 (2008), 309–322.
- [14] Shu Guo, Quan Wang, Lihong Wang, Bin Wang, and Li Guo. 2016. Jointly Embedding Knowledge Graphs and Logical Rules. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. 192–202.
- [15] Junheng Hao, Muhao Chen, Wenchao Yu, Yizhou Sun, and Wei Wang. 2019. Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1709–1719.
- [16] Mihály Héder. 2014. *Semantic Web for the Working Ontologist, Second edition: Effective modeling in RDFS and OWL*. *Knowl. Eng. Rev.* 29, 3 (2014), 404–405.
- [17] Robert Hoehndorf, Michel Dumontier, Anika Oellrich, Sarala Wimalaratne, Dietrich Rebholz-Schuhmann, Paul Schofield, and Georgios V. Gkoutos. 2011. A Common Layer of Interoperability for Biomedical Ontologies Based on OWL EL. *Bioinformatics* 27, 7 (April 2011), 1001–1008.
- [18] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard de Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. 2021. Knowledge graphs. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–37.
- [19] Yevgeny Kazakov, Markus Krötzsch, and František Simančík. 2014. The Incredible ELK. *Journal of Automated Reasoning* 53, 1 (2014), 1–61.
- [20] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations*.
- [21] Markus Krötzsch. 2012. OWL 2 Profiles: An Introduction to Lightweight Ontology Languages. In *Reasoning Web. Semantic Technologies for Advanced Query Answering. Reasoning Web 2012*. Vol. 7487. 112–183.
- [22] Maxat Kulmanov, Wang Liu-Wei, Yuan Yan, and Robert Hoehndorf. 2019. EL Embeddings: Geometric Construction of Models for the Description Logic \mathcal{EL}^{++} . In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. 6103–6109.
- [23] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. 2181–2187.
- [24] Julian Mendez. 2012. Jcel: A Modular Rule-Based Reasoner. In *Proceedings of the 1st International Workshop on OWL Reasoner Evaluation*, Vol. 858.
- [25] Sutapa Mondal, Sumit Bhatia, and Raghava Mutharaju. 2021. EmEL++: Embeddings for \mathcal{EL}^{++} Description Logic. In *Proceedings of the AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering (CEUR Workshop Proceedings, Vol. 2846)*.
- [26] Christopher J. Mungall, Carlo Torniai, Georgios V. Gkoutos, Suzanna E. Lewis, and Melissa A. Haendel. 2012. Uberon, an Integrative Multi-Species Anatomy Ontology. *Genome Biology* 13, 1 (2012), R5.
- [27] Mojtaba Nayyeri, Chengjin Xu, Mirza Mohtashim Alam, Jens Lehmann, and Hamed Shariat Yazdi. 2021. LogicENN: A Neural Based Knowledge Graphs Embedding Model with Logical Rules. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [28] Mojtaba Nayyeri, Chengjin Xu, Sahar Vahdati, Nadezhda Vassilyeva, Emanuel Sallinger, Hamed Shariat Yazdi, and Jens Lehmann. 2020. Fantastic Knowledge Graph Embeddings and How to Find the Right Space for Them. In *The Semantic Web*. 438–455.
- [29] Özgür Lütfü Özçep, Mena Leemhuis, and Diedrich Wolter. 2020. Cone Semantics for Logics with Negation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, Vol. 2. 1820–1826.
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, Vol. 32. 8024–8035.
- [31] Xi Peng, Zhenwei Tang, Maxat Kulmanov, Kexin Niu, and Robert Hoehndorf. 2022. Description Logic \mathcal{EL}^{++} Embeddings with Intersectional Closure. [arXiv:2202.14018](https://arxiv.org/abs/2202.14018)
- [32] Alan L Rector, Jeremy E Rogers, and Pam Pole. 1996. The GALEN high level ontology. In *Medical Informatics Europe '96*. IOS Press, 174–178.
- [33] Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. 2015. Injecting Logical Background Knowledge into Embeddings for Relation Extraction. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1119–1129.
- [34] Stefan Schulz, Boontawe Sunitisrivaraporn, Franz Baader, and Martin Boeker. 2009. SNOMED Reaching Its Adolescence: Ontologists’ and Logicians’ Health Check. *International Journal of Medical Informatics* 78 (April 2009), S86–S94.
- [35] Fatima Zohra Smaili, Xin Gao, and Robert Hoehndorf. 2019. OPA2Vec: Combining formal and informal content of biomedical ontologies to improve similarity-based prediction. *Bioinformatics* 35, 12 (2019), 2133–2140.
- [36] Barry Smith, Michael Ashburner, Cornelius Rosse, Jonathan Bard, William Bug, Werner Ceusters, Louis J Goldberg, Karen Eilbeck, Amelia Ireland, Christopher J Mungall, et al. 2007. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature biotechnology* 25, 11 (2007), 1251–1255.
- [37] Steffen Staab and Rudi Studer. 2010. *Handbook on ontologies*. Springer Science & Business Media.
- [38] Damian Szklarczyk, Annika L Gable, Katerina C Nastou, David Lyon, Rebecca Kirsch, Sampo Pyysalo, Nadezhda T Doncheva, Marc Legeay, Tao Fang, Peer Bork, Lars J Jensen, and Christian von Mering. 2021. The STRING Database in 2021: Customizable Protein–Protein Networks, and Functional Characterization of User-Uploaded Gene/Measurement Sets. *Nucleic Acids Research* 49, D1 (2021), D605–D612.
- [39] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2016. Complex Embeddings for Simple Link Prediction. In *Proceedings of The 33rd International Conference on Machine Learning*. PMLR, 2071–2080.
- [40] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (2017), 2724–2743.
- [41] Quan Wang, Bin Wang, and Li Guo. 2015. Knowledge Base Completion Using Embeddings and Rules. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [42] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 28.
- [43] Yuejia Xiang, Ziheng Zhang, Jiaoyan Chen, Xi Chen, Zhenxi Lin, and Yefeng Zheng. 2021. OntoEA: Ontology-guided Entity Alignment via Joint Knowledge Graph Embedding. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. 1117–1128.
- [44] Bo Xiong, Nico Potyka, Trung-Kien Tran, Mojtaba Nayyeri, and Steffen Staab. 2022. Faithful Embeddings for \mathcal{EL}^{++} Knowledge Bases. In *Proceedings of the 21st International Semantic Web Conference*, Vol. 13489. 22–38.
- [45] Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *Proceedings of the International Conference on Learning Representations*.
- [46] Xiaofang Zhai, Lei Huang, and Zhifeng Xiao. 2010. Geo-spatial query based on extended SPARQL. In *2010 18th International Conference on Geoinformatics*. 1–4.