

Ontology Embedding: A Survey of Methods, Applications and Resources

Jiaoyan Chen , Olga Mashkova , Fernando Zhapa-Camacho , Robert Hoehndorf , Yuan He ,
and Ian Horrocks 

(Survey Paper)

Abstract—Ontologies are widely used for representing domain knowledge and meta data, playing an increasingly important role in Information Systems, the Semantic Web, Bioinformatics and many other domains. However, logical reasoning that ontologies can directly support are quite limited in learning, approximation and prediction. One straightforward solution is to integrate statistical analysis and machine learning. To this end, automatically learning vector representation for knowledge of an ontology i.e., *ontology embedding* has been widely investigated. Numerous papers have been published on ontology embedding, but a lack of systematic reviews hinders researchers from gaining a comprehensive understanding of this field. To bridge this gap, we write this survey paper, which first introduces different kinds of semantics of ontologies and formally defines ontology embedding as well as its property of faithfulness. Based on this, it systematically categorizes and analyses a relatively complete set of over 80 papers, according to the ontologies they aim at and their technical solutions including geometric modeling, sequence modeling and graph propagation. This survey also introduces the applications of ontology embedding in ontology engineering, machine learning augmentation and life sciences, presents a new library mOWL and discusses the challenges and future directions.

Index Terms—Ontology, ontology embedding, web ontology language, knowledge graph, representation learning.

I. INTRODUCTION

ONTOLOGIES are formal, explicit and shared representations of knowledge within a domain, with definitions and axioms for concepts, properties, relations and other types

Received 16 June 2024; revised 13 February 2025; accepted 5 April 2025. Date of publication 11 April 2025; date of current version 28 May 2025. This work was supported in part by EPSRC projects OntoEm under Grant EP/Y017706/1, in part by ConCur under Grant EP/V050869/1, in part by U.K. FIRES under Grant EP/S019111/1, in part by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Grant URF/1/4675-01-01, Grant URF/1/4697-01-01, Grant URF/1/5041-01-01, Grant REI/1/5659-01-01, Grant REI/1/5235-01-01, and Grant FCC/1/1976-46-01, and in part by the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence (SDAIA-KAUST AI). Recommended for acceptance by F. Chiang. (Olga Mashkova and Fernando Zhapa-Camacho contributed equally to this work.) (Corresponding author: Jiaoyan Chen.)

Jiaoyan Chen is with the University of Manchester, M13 9PL Manchester, U.K., and also with the University of Oxford, OX1 2JD Oxford, U.K (e-mail: jiaoyan.chen@manchester.ac.uk).

Olga Mashkova, Fernando Zhapa-Camacho, and Robert Hoehndorf are with the King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia.

Yuan He and Ian Horrocks are with the University of Oxford, OX1 2JD Oxford, U.K.

Digital Object Identifier 10.1109/TKDE.2025.3559023

of entities [1]. Ontologies have been a critical technology in Knowledge Management, Information Systems, the Semantic Web, Natural Language Processing and Artificial Intelligence, playing an increasingly important role in many fields such as Healthcare, Bioinformatics and E-commerce. A simple ontology can just be a set of concepts arranged in a hierarchy with the subsumption (inclusion) relationship between concepts indicating that instances of one concept all belong to another. Such ontologies are capable of representing taxonomies of domains, such as the BBC Wildlife Ontology [2] and the International Classification of Diseases (ICD) [3], dating back to Porphyrian Tree for presenting Aristotle’s categories in the third century AD. Meanwhile, many data and knowledge management systems such as e-commerce platforms [4] also adopt such simple ontologies for type information of data.

With the fast development of the Web in 1990 s, representing and exchanging data and knowledge on the Web became desirable. To this end, several standards were proposed for defining more complex ontologies for constructing the Semantic Web [5], [6]. In 1999, Resource Description Framework (RDF)¹ which defines the syntax of triple (*Subject, Predicate, Object*) was proposed for representing data, and in 2000, the vocabulary of RDF Schema (RDFS)² was proposed for building ontologies as the schemas of data [7]. The vocabulary of RDFS can represent not only hierarchical concepts, but also instance membership to concepts, property hierarchies, and property domains and ranges.

The Web Ontology Language (OWL), including its second version OWL 2, was published upon the foundation of RDF and RDFS for building ontologies that can represent more complex knowledge with logics such as the disjunction, conjunction and disjointness of concepts, and the existential and universal rules.³ OWL was underpinned by Description Logic — a fragment of first-order logic with decidable reasoning and efficient decision procedures [8]. Many widely used ontologies, such as the Gene Ontology (GO) [9], the Food Ontology (FoodOn) [10], the DBpedia ontology [11] and the aforementioned BBC Wildlife Ontology, adopt OWL [12].

Besides the formally and explicitly defined semantics, most real-world ontologies are also equipped with literals, including attribute values of instances defined by data properties,

¹<https://www.w3.org/TR/rdf11-concepts/>

²<https://www.w3.org/TR/rdf-schema/>

³<https://www.w3.org/TR/owl-features/>

and meta-data associated with entities⁴ defined by annotation properties,⁵ representing information of name, definition, comment, image, version and so on. For example, the concept *obo:FOODON_00002873* in FoodOn is associated with an English name “okara”, a synonym “soy pulp”, a long definition “Okara, soy pulp, or tofu dregs is a pulp consisting of insoluble parts of the soybean that...”, the source of this definition, and an image of okara [10]. These annotations, originally created for human understanding, contain important information that is often complementary to the formal semantics. However, they cannot be utilised by symbolic reasoning.

In early 2010 s, word embedding algorithms like Word2Vec were proposed to represent natural language words as low dimensional vectors, capturing their semantic relationships like correlations within the vector space [13]. Similar ideas of representation learning were also applied to knowledge graphs (KGs) that are composed of relational facts, giving rise to some classic algorithms like TransE [14] and RDF2Vec [15]. The instances and relations (i.e., object properties) are embedded with their semantics indicated by facts retained in the vector space. Take TransE as an example; the embedding of a KG, denoted as a mapping function $v(\cdot)$ from its instances and relations to vectors, is learned such that each relational fact (h, r, t) where the head instance h , the relation r and the tail instance t correspond to the subject, predicate and object in an RDF triple, respectively, is kept in the vector space as $v(h) + v(r) \approx v(t)$. Such embeddings not only enable machine learning and statistical algorithms to utilise the knowledge, but also support neural-symbolic reasoning within a KG, with both approximation and prediction [16], [17].

Similarly, representation learning has also been applied to ontologies for embedding. Some early works such as [18] and [19] proposed to embed triples with relations of transitivity, including those from WordNet [20]. Their methods are applicable to simple ontologies with concept hierarchies. However, the semantics of RDFS and OWL ontologies are much more complex, and, accordingly, more advanced vector representation models, which use high dimensional balls, boxes and axis-aligned cones for representing concepts, were recently proposed, following some early works including EmbedS [21] for RDFS ontologies, EEmbeddings [22] and E2R [23] for OWL ontologies of Description Logic \mathcal{EL}^{++} and \mathcal{ALC} , respectively. Meanwhile, several complex embedding frameworks such as OPA2Vec [24] and OWL2Vec* [25] were also proposed to embed both formal semantics and textual literals, often upon sequence learning methods.

Such ontology embedding methods have brought new solutions for ontology engineering and dramatically extended the application of ontologies. Many of them have been verified in tasks with real-world data, including concept subsumption inference [25], the domain application of protein–protein interaction prediction [24], and ontology augmented zero-shot and few-shot learning [26], [27]. Besides ontology embedding, there are some

other directions that attempt to combine ontology with machine learning or statistical approach, including fuzzy ontology standards such as Fuzzy OWL [28], traditional ontology learning methods such as the inductive logic programming system DL-Leaner [29], [30], and some neural-symbolic frameworks such as Logic Tensor Network [31]. In comparison with them, ontology embedding methods focus on the automatic learning of the given knowledge’s vector representations so as to supporting the integration with different machine learning and statistical models.

Briefly, there are quite a few results about ontology embeddings, covering theoretic analysis, new methods and applications. Although some survey papers for KG embedding involve ontologies [32], [33], [34], but they only analyse those embedding methods that aim at relational facts, regarding simple ontologies as a kind of additional constraints. There is a shortage of systematic review to papers on ontology embeddings. Kuzmanov et al. [35] reviewed some works in 2021 on using machine learning for analyzing semantic similarity with ontologies. Several ontology embedding methods including Onto2Vec [36], OPA2Vec [24] and EL Embeddings [22] are covered, but they are far from complete, especially considering there are quite a few papers published after 2021.

This survey aims to bridge the above gap, with (i) systematic categorization and comparison of the ontology embedding methods according to the employed techniques and the target ontologies, (ii) review of the applications in supporting knowledge engineering, machine learning and life science knowledge discovery together with benchmarks and metrics, (iii) introduction and result demonstration of a library named mOWL [37] that can support the implementation of ontology embedding methods, and (iv) discussion on the challenges and future directions. This survey has reviewed over 80 papers (around 40 of them are for new embedding methods) published in conferences and journals of Computer Science, AI and Bioinformatics, covering all the relevant works on ontology embedding, to the best of our knowledge. We believe it will benefit all the researchers who are interested in some topics among ontology, KG, knowledge representation, semantic embedding, semantic techniques, knowledge engineering, neural-symbolic integration, bioinformatics, and AI for life sciences.

The remainder of this paper is organized as follows. Section II gives the background of ontology and semantic embedding. Section III reviews ontology embedding methods. Sections IV and V review the applications. Section VI demonstrates mOWL. Section VII presents our perspectives on challenges and future directions. Section VIII concludes the paper.

II. BACKGROUND

A. Symbolic Knowledge Representation With Ontologies

1) *Knowledge Graph (KG)*: In this paper, we distinguish KGs from ontologies. We refer to KGs as those knowledge bases mainly composed of relational facts in RDF, following the definition in most KG embedding papers [16], [38]. A KG can be denoted as $\mathcal{G} = \{\mathbf{I}, \mathbf{R}, \mathbf{T}\}$. \mathbf{I} denotes a set of instances (also known as entities), corresponding individuals in Description Logic. \mathbf{R} denotes a set of binary relations. \mathbf{T} denotes a set of relational facts, i.e., $\mathbf{T} = \{(h, r, t) | h, t \in \mathbf{I}, r \in \mathbf{R}\}$, where

⁴ In the community of ontology, a concept is modeled as a class, an instance corresponds to an individual in Description Logic, and the term “entity” includes class, instance and property. In the community of knowledge graph, the mention of an entity is actually equivalent to an instance. For clarity, we adopt the terms of the ontology community in this paper.

⁵ <https://www.w3.org/TR/owl-ref/#Annotations>

(h, r, t) is an RDF triple, h and t , as the subject and object, are also called as the triple's head and tail, respectively. One simple example is $(Bob, hasFather, Alex)$. Sometimes (h, r, t) is also denoted in form of a relation $r(h, t)$. For real-world KGs of the Semantic Web, each instance or relation should be uniquely identified, usually by an Internationalized Resource Identifier (IRI).

2) *RDF Schema (RDFS)*: RDFS can define either an ontological schema for a KG or an independent ontology with the following main features.

- RDFS can define a set of concepts (classes) C and assert the membership of instances using the built-in predicate $rdfs:type$. For example, the triple $(Alex, rdfs:type, Father)$ represents that $Alex$ is an instance of $Father$.
- RDFS can define the subsumption between concepts with the built-in predicate $rdfs:subClassOf$. Considering two concepts $Father$ and $Parent$ defined in C . The triple $(Father, rdfs:subClassOf, Parent)$ represents that $Parent$ subsumes $Father$ and $Father$ is a sub-concept of $Parent$.
- RDFS can define the domain and range of a relation (i.e., object property) with the built-in predicates of $rdfs:domain$ and $rdfs:range$, respectively. If the domain (resp. range) of a relation r is a concept C , the heads (resp. tails) that are associated with r must be declared or inferred to be instances of C . For example, we can define the range of $hasFather$ to $Father, Parent$ and/or $Male$. RDFS can also define the range of a data property with built-in data types.
- RDFS can define the subsumption between two properties with the built-in predicate $rdfs:subPropertyOf$, indicating that the instance pairs associated to one property all belong to those associated to another property. One example is $(hasFather, rdfs:subPropertyOf, hasParent)$.

3) *Description Logic (DL) and Web Ontology Language (OWL)*: OWL has different sub-languages and comes in multiple versions. The complete languages, OWL Full and OWL 2 Full, are not decidable. In this part we mainly introduce (i) the vocabularies of OWL 2 DL which are defined based on the DL fragment \mathcal{SROIQ} , and (ii) some of its widely used sub-languages that are developed for different scenarios with a better balance between knowledge expressivity and reasoning complexity. A signature consists of three finite sets of symbols: a set N_I of individual names, a set N_C of concept names, and a set N_R of role names. DL \mathcal{SROIQ} allows recursive concept definition as

$$\begin{aligned} & \top \mid \perp \mid A \mid C \sqcap D \mid C \sqcup D \mid \neg C \mid \exists r.C \mid \forall r.C \mid \\ & \geq nr.C \mid \leq nr.C \mid \exists r.Self \mid \{a\} \end{aligned}$$

where \top is the top concept, \perp is the bottom concept, $A \in N_C$ is an atomic (or named) concept, $r \in N_R$ is an atomic role (equivalent to a binary relation), $a \in N_I$ is an individual, C and D are themselves (possibly complex) concepts, n is a number of cardinality, $\exists r.Self$ is a concept that indicates the set of entities in the domain that are related by r with themselves [8]. We say a concept is *complex* when it is constructed with one or multiple logical operators such as $\sqcap, \sqcup, \exists, \forall$ and \neg . A DL ontology \mathcal{O} can be composed of a TBox \mathcal{T} and an ABox \mathcal{A} . The TBox defines logical background knowledge in the form of concept

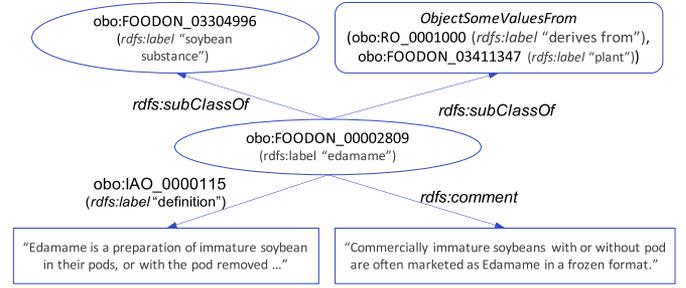


Fig. 1. A fragment from the OWL ontology FoodOn [10].

subsumption axioms $C \sqsubseteq D$ (Generalized Concept Inclusion, GCI), and role axioms for logical background knowledge of role composition, role subsumption, and role characteristics like functionality, transitivity and so on. Sometimes these role axioms are separately divided into an RBox, and accordingly the DL ontology is composed of a TBox, an RBox and an ABox. The ABox contains concrete data including concept assertions in form of $C(a)$, and role assertions in form of $r(a, b)$. With the defined logic, symbolic reasoners can be applied to infer hidden knowledge (i.e., entailment reasoning), check the ontology consistency and find justification that leads to inconsistency.

\mathcal{ALC} and \mathcal{EL}^{++} are two important fragments of DL that are widely investigated in ontology embedding. \mathcal{ALC} is known as *Attributive Concept Language with Complements* and allows recursive concept definition with $\top \mid \perp \mid A \mid C \sqcap D \mid C \sqcup D \mid \neg C \mid \exists r.C \mid \forall r.C$ [8]; \mathcal{ALC} is a prototypical DL mainly studied because it has most of the major features.

The DL fragment \mathcal{EL}^{++} which corresponds to OWL 2 EL profile allows recursive concept definition with $\top \mid \perp \mid A \mid C \sqcap D \mid \exists r.C \mid \{a\}$ [39]. Due to a high knowledge representation capability but a polynomial time complexity in reasoning, DL \mathcal{EL}^{++} is widely used by many real-world ontologies such as SNOMED CT [40]. Example 1 presents a toy family ontology of DL \mathcal{EL}^{++} , with a TBox and an ABox.

Example 1: The following \mathcal{EL}^{++} ontology models a simple family domain:

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

$$\mathcal{A} = \{ Father(Alex), Child(Bob), hasParent(Bob, Alex) \}$$

The ontology in Example 1 can be implemented with the vocabularies defined in the standards of RDF, RDFS and OWL; for example, $rdfs:subClassOf$ for concept subsumption, $owl:ObjectSomeValuesFrom$ for the existential quantification $\exists r.C$ and $owl:ObjectAllValuesFrom$ for the universal quantification $\forall r.C$. Fig. 1 presents a fragment for the concept $obo:FoodON_00002809$ (“edamame”) which is a sub-concept of a named concept $obo:FOODON_03304996$ (“soybean substance”) and an existential quantification with the property of $obo:RO_0001000$ (“derives from”) and the concept of $obo:FOODON_03411347$ (“plant”).

4) *Ontology Literals*: The literals of an ontology are mainly defined in two approaches. (1) The first approach is to associate instances with literals by datatype properties, which are sometimes known as attributes and whose values can be of different types such as natural language phrases and long text, real values, data and time, category, image and domain specific sequence (e.g., gene sequence). For example, instances of *Person* may have address, height, birth data and so on. (2) The second approach is to associate entities with meta information by annotation properties. In real-world ontologies, most such literals are in form of text. For example, in the ontology fragment in Fig. 1, the concepts and properties are annotated with names by the built-in vocabulary of *rdfs:label*,⁶ and *obo:FOODON_00002809* is annotated with two long sentences — a comment by *rdfs:comment* and a definition by the ad-hoc annotation property *obo:IAO_0000115*. Some other literals such as links to external sources, editors and images are also widely used. Although all these literals include important information, as informal semantics, they are often ambiguous and cannot be used for inference by symbolic ontology reasoners. *In ontology embedding, ontologies either with or without literals are considered.*

5) *Target Ontologies*: In our context of ontology embedding, “ontology” refers either to a TBox, or any DL ontology with a non-empty TBox, i.e., $(TBox, ABox, RBox)$ or $(TBox, ABox)$ with $TBox \neq \emptyset$. Different embedding methods aim at different semantics of an ontology. Accordingly, we divide the ontologies considered in the current ontology embedding works into four kinds:

- *Simple Ontology* which refers to DL ontologies that have only a TBox composed of the top concept \top (e.g., *owl:Thing* in OWL ontologies), named concepts and subsumption axioms between named concepts. They are equivalent to taxonomies composed of hierarchical classes.
- *Complex Ontology* which refers to DL ontologies that have a TBox containing any concept definitions of DL *SRIOQ* beyond \top . These ontologies may also contain an RBox, an ABox, or both.
- *Ontology with Literals* which refers to simple or complex ontologies that have literals.
- *Ontology with KG* which refers to KGs composed of large scale relational facts, equipped with DL ontologies that has a TBox (and an RBox in option) as the KG’s schema.

B. Semantic Embedding

In this part, we first summarise word embedding and KG embedding, and then give formal definitions and properties of general embedding and ontology embedding.

1) *Word Embedding*: Word embedding algorithms such as Word2Vec [13] and GloVe [41] learn vector representations of tokens (which are either words or sub-words) from a large corpus concerning their semantics such as co-occurrence in the sentences. Taking Word2Vec as an example, it learns a Feed-Forward Neural Network from natural language sentences by one of the two auto-encoding architectures — continuous skip-gram which predicts the surrounding tokens of each token

⁶ The textual value often has a language tag and is regarded as English by default. The entity IRIs sometimes also indicate the name information.

and continuous Bag-of-Words which predicts a token based on its surrounding tokens. For each token, the hidden layer output of the network is its embedding. Tokens with more similar meanings are expected to have higher vector similarities. Such embedded semantics can partially support analogical reasoning, e.g., $v(\textit{king}) - v(\textit{father}) \approx v(\textit{queen}) - v(\textit{mother})$.

Embeddings by Word2Vec and GloVe are *non-contextual*, which means each token has one unique vector representation no matter what surrounding tokens it has. Recently, with the development of Transformer-based encoder architectures [42] and Pre-trained Language Models (PLMs) like BERT [43], *contextual word embeddings* have been widely developed and adopted. Taking BERT as an example, it learns a Transformer architecture from a corpus by predicting the masked token in each sentence and the next sentence of a given sentence. The vector of a token is based on the attentions from itself and its surrounding tokens in the sentence. Given a sentence “the bank robber was seen on the river bank”, the first “bank” and the second “bank” have different vectors due to different surrounding tokens. Such contextual word embeddings encode more semantics and often perform better than non-contextual word embeddings in many tasks.

Besides natural language text, the above contextual and non-contextual word embedding techniques are also applicable to other kinds of sequential data, such as BioVec for biological sequences like genes and proteins [44], and Node2Vec for paths from graphs [45]. Their great success in many domains also motivate researchers to applying them to KGs and ontologies. For simplicity, in this paper we call word embedding as well as representation learning for other sequential data as sequence learning.

2) *Knowledge Graph Embedding*: These methods learn vector representations of instances and relations from the relational facts in a KG with their semantics retained in the vector space [16]. In general, typical technical solutions can be divided into *geometric modeling* such as the translational method TransE [14], *tensor decomposition* such as the bilinear method DistMult [46], *neural network modeling* such as ConvE [47], *random walk-based sequence modeling* such as RDF2Vec [15]. TransE has been introduced in Section I. Taking RDF2Vec as another example, it first conducts random walk for extracting paths composed of instances and relations, and then learns a Word2Vec model for encoding them. KG embedding has also been extended to encode semantics beyond relational facts, especially textual literals in combination with word embedding methods [48], [49], and logics such as horn rules, schemas and constraints with additional modeling methods [32], [34].

3) *Formal Definitions of Embedding*: Although the term “embedding” has been widely used in contexts of machine learning, NLP and KG, we find it useful to explicitly make our understanding of embedding clear here as it will provide a guide for analyzing and classifying ontology embedding methods. We start with the definition of “embedding” as it is used in mathematics:

Definition 1 (Embedding (mathematics)): An embedding is an injective and structure-preserving map between two mathematical structures which can be algebraic, topological, or geometrical structures.

In machine learning, embedding is used in a somewhat different sense:

Definition 2 (Embedding (machine learning)): An embedding e is a learned mapping between the elements of a mathematical structure and the elements of some structure \mathcal{S} .

The notion of “embedding” in machine learning is therefore related to, but less strict than, a structure-preserving map, and embeddings are representations learned from data, largely built upon the foundation of representation learning [50]. The learned representations are usually not arbitrary but rather aim to preserve some “semantics” of the original data. We capture this by defining the notion of “faithfulness”:

Definition 3 (Faithfulness of embedding (machine learning)): An embedding (machine learning) e is “faithful” if e converges to some embedding (mathematics) f .

4) **Definitions and Properties of Ontology Embedding:** An ontology has a syntactic structure which consists of a set of logical symbols like connectives and quantifiers, a set of non-logical symbols including constants, functions and relations, a set of well-formed formulas constructed from the logical and non-logical symbols, a set of axioms (a subset of the well-formed formulas), and a set of inference rules for deriving new formulas from the axioms and previously derived formulas. The signature of the ontology includes concept, individual and relation symbols. The ontology’s model structure over its signature consists of a non-empty set D , called the domain or universe of the model, an interpretation function I that assigns each individual symbol to an element of D , each concept symbol to a subset of D , and each binary relation symbol to a binary relation on D (i.e., a subset of D^2). The semantic structure of the ontology \mathcal{O} , denoted as $Mod(\mathcal{O})$, consists of the collection of all model structures that satisfy the axioms in \mathcal{O} . $Mod(\mathcal{O})$ can be seen as a class or a category in the sense of category theory, depending on the context. With these background, we formally define ontology embedding and its faithfulness:

Definition 4 (Ontology embedding): Let \mathcal{O} be an ontology with signature $\Sigma_{\mathcal{O}}$. An ontology embedding is an embedding (machine learning) e of the term algebra $T(\Sigma_{\mathcal{O}})$ generated by the signature $\Sigma_{\mathcal{O}}$ and the well-formedness rules of the underlying Description Logic for constructing concept descriptions, role expressions, and axioms.

Definition 5 (Faithful ontology embedding): An embedding of an ontology \mathcal{O} is *faithful* if it converges to an embedding (mathematics) which preserves some mathematical structure \mathcal{S} of \mathcal{O} .

Here, we only consider embeddings for specific ontologies (with a specific signature), not embeddings of an entire logic. The notion of “faithfulness” for embeddings requires specifying a mathematical structure \mathcal{S} to which the embeddings are faithful. There are at least three mathematical structures that can be assigned to ontologies: (1) the syntactic structure, potentially combined with syntactic inference rules (i.e., the deductive calculus); (2) a single arbitrary model of the ontology, i.e., a model structure in which all ontology axioms are true; and (3) the semantic structure of the ontology [51]. The mathematical structure that the embedding aims to preserve can also be used to classify and distinguish ontology embedding models.

Another property of ontology embedding is “interpretability”. In machine learning, there is no universally-agreed definition of interpretability [52]; in [53], interpretability is described as “the ability to explain or to present in understandable terms

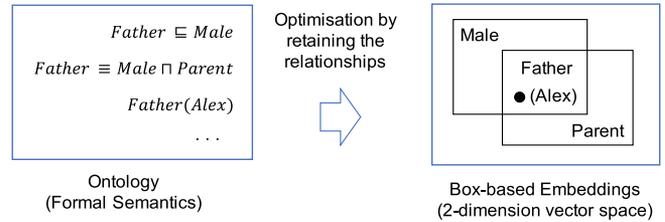


Fig. 2. Demonstration of the ontology embedding solution of geometric modeling with the examples of ELBE and Box²EL.

to a human”. In the context of ontology embedding, we can strictly define “interpretability” as the ability to reconstruct both symbols and composition rules from their embeddings. Faithful ontology embeddings are interpretable in that sense since injectivity ensures that each symbol can be uniquely recovered from its embedded representation, and the preservation of a mathematical structure of an ontology inherent in faithful embeddings allows for the restoration of the composition rules between symbols. A non-faithful ontology embedding does not support full restoration, but may still allow restoring a part of the symbolic semantics and justifying the inference with the embeddings. This characteristic is important for assessing ontology embedding methods since it reflects to what extent humans, who inherently engage in symbolic reasoning, can read off ontology axioms and understand the reasoning from the constructed embedding.

III. ONTOLOGY EMBEDDING METHODS

In this section, we first analyse the main technical solutions that have been commonly adopted for ontology embedding (Section III-A), and then review the ontology embedding works for each kind of ontologies (Sections III-B to III-E).

A. General Technical Solutions

The technical solutions that are commonly adopted by the current ontology embedding methods include the following:

- **Geometric Modeling** aims to generate embeddings that are faithful to some model structures of an ontology; they generate (or approximate) logical models of an ontology by interpreting concepts as geometric regions, interpret individuals as members of these regions, and relations as pairs of points standing in some geometric relation. Take the embedding algorithms ELBE [54] and Box²EL [55] for the DL \mathcal{EL}^{++} as an example. They model an instance as a high dimensional point, i.e., one single vector, and models a concept as a high dimensional axis-aligned box represented by one vector for the box center and one vector for the box offset. The instantiation relation between an instance and a concept is modeled as the instance’s point lying within the concept’s box. The basic idea is demonstrated in Fig. 2 with a toy example. These types of embeddings are highly interpretable since they induce geometric relations between geometric objects representing relations, concepts and individuals which align with ontology axioms.

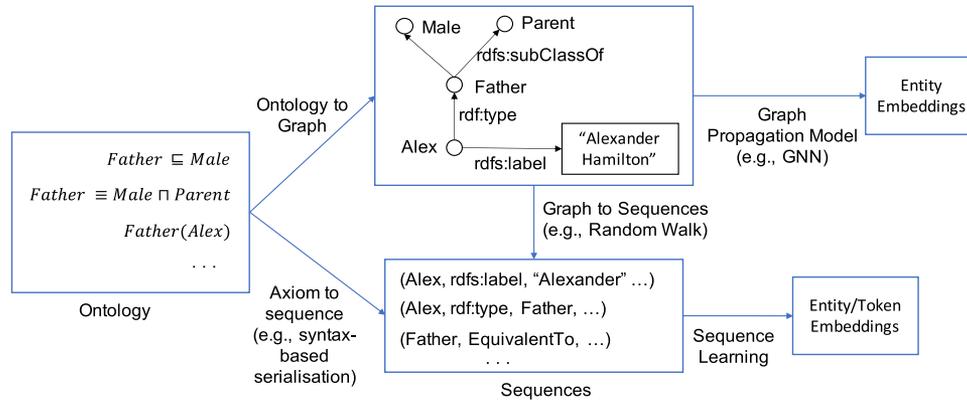


Fig. 3. Demonstration of the ontology embedding solutions of sequence modeling and graph propagation.

TABLE I
IMPORTANT CHARACTERISTICS (PROS AND CONS) OF THE THREE TECHNICAL SOLUTIONS, AND THE ONTOLOGIES THAT MAIN METHODS OF EACH TECHNICAL SOLUTION AIM AT. ONTO. AND LIT. ARE SHORT FOR ONTOLOGY AND LITERAL, RESPECTIVELY

Solution	Pros	Cons	Targeted Ontologies and Citations
Geometric Modeling	High interpretability; high faithfulness to the semantic structure	Hard to integrate literals; not support many features of OWL	Simple Onto.: [57][58][19][59][60][61][18] [62][63][64][65] Complex Onto.: [66][22][23][67][68] [54][69][70][71][72][55] Onto. & Lit.: [26] Onto. & KG: [21][73][74][75][76][77][78]
Sequence Modeling	Extensible to different ontologies; support literals; learn correlations as effective features for downstream tasks	Low interpretability; only partial faithful to correlations and co-occurrences	Complex Onto.: [79][36][80] Onto. & Lit.: [81][82][83][24][84] [85] [25][86][87][88][89][90][91]
Graph Propagation	Well embed the graph especially the concept hierarchy	Low interpretability; only partially faithful to the graph structure	Simple Onto.: [92] Onto. & Lit.: [93][56]

- *Sequence Modeling* first transforms the ontology axioms and literals into sequences composed of entities (and tokens), and then adopts a sequence learning model to learn their embeddings. The basic idea is demonstrated in Fig. 3. Take OWL2Vec* [25] as an example. It first extracts sequences composed of entities and textual literals from an OWL ontology by mapping the ontology to a graph and performing multiple random walks over the graph, and then learns a word embedding model from the sequences for encoding the entities and tokens in the text. These methods are usually partially faithful to correlations and co-occurrences of symbols in the axioms of the ontology, and thus they have low interpretability.
- *Graph Propagation* represents an ontology by a (multi-relation) graph with initial node representations, and then learns a graph propagation model for new node representations. The basic idea is also demonstrated in Fig. 3. For example, for concept matching, the work [56] uses two Graph Convolutional Networks to learn embeddings of two cross-ontology concepts, respectively, by propagating the initial word embeddings of their surrounding concepts. Methods of this type are commonly interpretable to a limited degree since they preserve mainly the concept hierarchy.

Table I gives a summary of these technical solutions, including an analysis of the pros and cons according to two dimensions — *faithfulness* with regard to different mathematical structures of an ontology, and *interpretability* which indicates

how transparency the reasoning in the vector space is. Table I also classifies different ontology embedding methods of each technical solution based on the ontologies that they can deal with (i.e., the *targeted ontologies*).

B. Embedding Simple Ontology

The concept hierarchies of an ontology imply its basic structural semantics. We refer to ontology embedding that solely consider concept hierarchies as simple ontology embedding as it omits more complex logical semantics. Fig. 4 presents two dimensions — Embedding Method and Embedding Space, their values and corresponding works of simple ontology embedding.

Geometric modeling is the most prevalent technique for simple ontology embedding. This includes box embedding in euclidean space, hyperbolic distance-based embedding in hyperbolic space, cone embedding for both geometries, and graph propagation in hyperbolic space.

In the context of euclidean space, a key approach is to devise a function that preserves the hierarchical order of entities. A typical work by [57] proposes using the reversed product order, essentially forming a euclidean cone where each entity’s embedding value is at least that of its parent entity. Building on this, [58] extends to encode entities using probability densities instead of points. The box embedding is another typical construction which simulates the hierarchical ordering with hyper-rectangles, where a child entity’s box is consumed within its parent entity’s box. Unlike cone embeddings, which are typically parameterized by their apex, box embeddings require two vectors representing

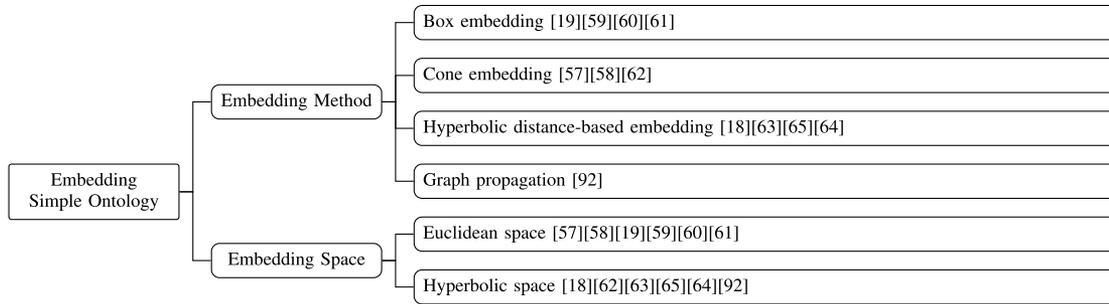


Fig. 4. Dimensions, their values and corresponding works of embedding simple ontology.

the minimal and maximal coordinates in the hyper-rectangle. To further improve box embeddings, [59] explores a probabilistic relaxation to achieve a smoother distribution, [60] investigates encoding dual hierarchical relationships (hypernym and meronym) at the same time, and [61] addresses the local identifiability issue, proposing the Gumbel-box process as a solution.

Hyperbolic space, with its expansive property and theoretical underpinning, is particularly suitable for representing hierarchical structures. The Poincaré embedding [18] is a typical approach that minimizes hyperbolic distances between related entities while maximizing the separation from unrelated ones in a unit Poincaré ball. This spatial arrangement places more general entities near the origin and more specific entities closer to the boundary, reflecting their hierarchical depth. However, numerical instabilities near the boundary of the manifold are a known challenge. To address this, [94] investigates an alternative hyperbolic model, the Lorenz (or Hyperboloid) model, while [65] proposes the extended Poincaré ball with geometric distortion. To augment the embedded semantics, [63] and [64] explores the integration of Poincaré embeddings with pre-trained word embeddings, while [92] utilises the hyperbolic graph convolutional networks (HGNC) [95] for aggregating neighbourhood information. In hyperbolic space, establishing geometric shapes like boxes and cones, which are straightforward in euclidean space, requires more nuanced considerations. A notable contribution in this line is the hyperbolic entailment cone by [62], which not only preserves transitivity but also enables direct prediction of entity subsumptions through its construction.

The techniques discussed here aim to achieve high geometric interpretability in encoding hierarchies, whether preserving order or exploiting geometric properties. Although not all cited works specifically test with ontology concept hierarchies, their methodologies are readily applicable to such structures.

C. Embedding Complex Ontology

Complex ontologies are mainly embedded by geometric modeling. The methods aim to map concepts, individuals and roles into a continuous vector space (e.g., \mathbb{R}^n) where they can be represented as points or geometrical regions. This allows to capture some aspects of semantics of the underlying ontology by means of geometric properties of the embedding space. We analyze complex ontology embedding methods from four perspectives: embedding method, semantics complexity, strategy to embed ABox and RBox axioms, and theoretical analysis; corresponding categorization and related works are shown in Fig. 5.

Multiple geometric models have been developed for the lightweight DL \mathcal{EL}^{++} . ELEmbeddings [22] and EmEL++ [68] represent named concepts as n -dimensional euclidean balls, which cannot faithfully model concept intersection since the intersection of two balls is no longer a ball in \mathbb{R}^n . To address this issue, boxes have been adopted [54], [55], [69]. These methods can be categorized further by the relation model implemented within their frameworks: ELBE [54] utilizes TransE [14]; this model based on vector translations cannot faithfully represent 1-to- N , N -to- N and N -to-1 relations. To overcome this limitation, Box²EL [55] adapts the relational model of BoxE [97], which also allows for capturing role composition and role subsumption. BoxEL [69] develops an alternative way to elucidate relations by affine transformations which solves the problem of concept embedding size incompatibility. In [71] a proper non-convex geometric interpretation for concepts, roles and individuals is introduced first through mapping of interpretation domain elements into binary vectors, and then the convex hulls of constructed regions are considered. This method's applicability remains an open question since it lacks implementation and empirical evaluation. [70] is another method that is mainly described from a theoretical perspective. It uses axis-aligned cones for concept interpretations in \mathcal{ALC} ontologies. As opposed to previously discussed works, this method builds partial models: as an example, if for some individual a only assertion axioms of $(C \sqcup D)(a)$ and $C(a)$ are presented within the ontology and neither $D(a)$ nor $(\neg D)(a)$ can be proven, the embedding leverages multiple interpretations since there are several dissimilar ways to interpret a within this geometric framework. There are some other methods tailored to encode \mathcal{ALC} theories, using fuzzy sets [96] or ordered vector spaces [72].

From the perspective of the expressivity of the DL embedded, many geometric models focus on encoding the constructs of \mathcal{EL}^{++} . ELEmbeddings [22], ELBE [54] and BoxEL [69] work with the \mathcal{EL} fragment enriched with nominals omitting role inclusions. Subsumption axioms are interpreted as the containment of one geometric region within another one. EmEL++ [68] and Box²EL [55] include objective functions for role inclusion and role chain axioms adopting geometrical containment for role interpretations: in Box²EL, a box inclusion loss is used towards boxes that represent the head and tail parts of a role for role inclusion and chain axioms; in EmEL++ [68], the role hierarchy is interpreted via establishing partial order for vectors. [71] provides detailed information about each element in the interpretation domain (including to what individual it corresponds, in which concepts it is contained and how it is related to other elements); while this approach can be applied to \mathcal{ALC} , it requires

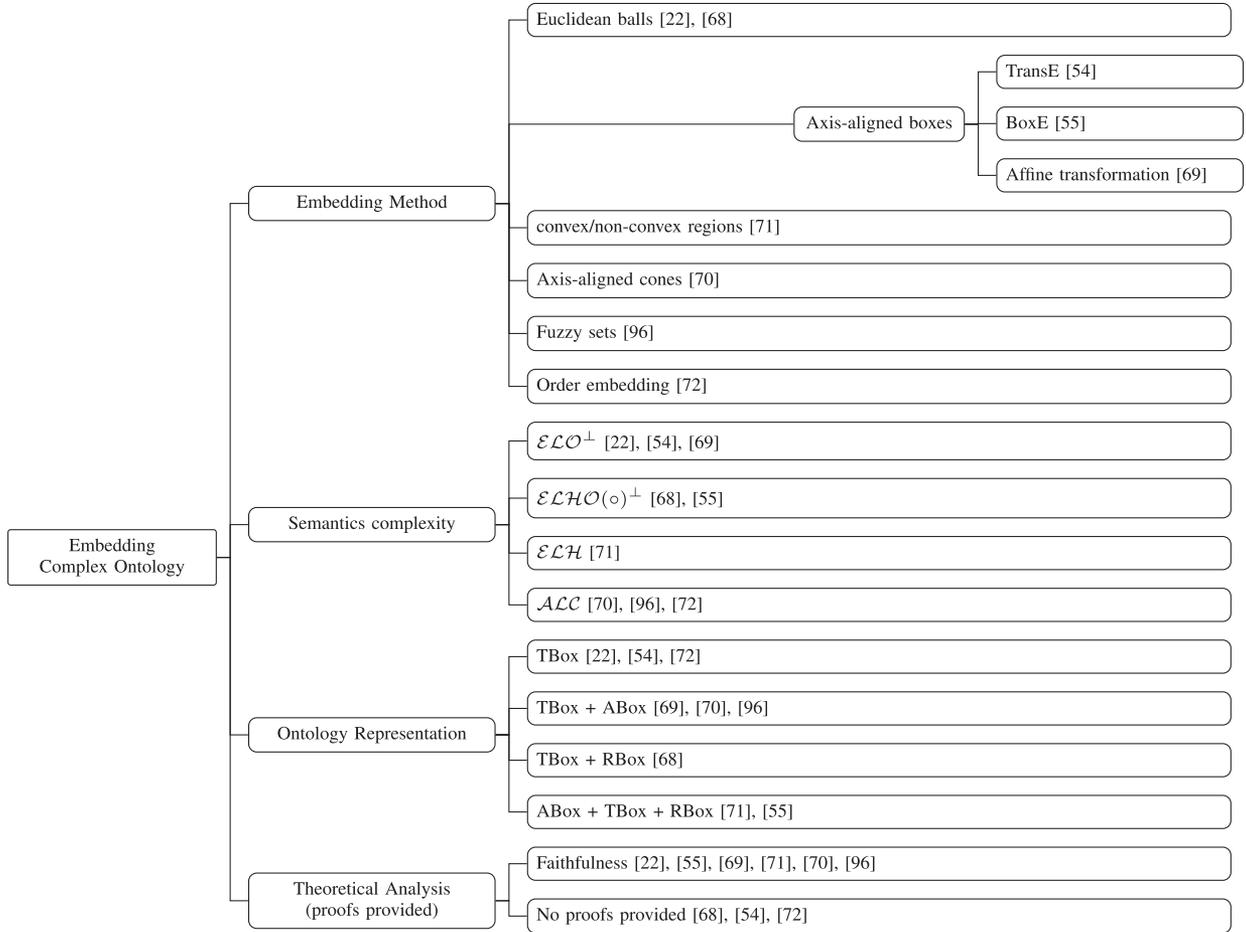


Fig. 5. Dimensions, their values and corresponding works of embedding complex ontology.

to store sparse vectors of size $|\mathcal{N}_I| + |\mathcal{N}_C| + |\mathcal{N}_R| \cdot |\Delta^{\mathcal{I}}|$, where $\Delta^{\mathcal{I}}$ is the interpretation domain. Other methods under consideration [70], [72], [96] target the full expressivity of \mathcal{ALC} ; in particular, the negation of a concept is incorporated as a membership function of $(\neg C)^{\mathcal{I}}$ assigned to all individuals via fuzzy negation operator applied to the corresponding membership function of $C^{\mathcal{I}}$ [96], by concept lattice enrichment [72], or using polar operators [70].

An ontology may have an ABox and an RBox, besides a TBox. Some embedding methods do not distinguish between individuals and concepts. They eliminate ABox and work solely with TBox axioms [22], [54], [72] by translating assertional axioms $C(a)$ and $r(a, b)$ into TBox axioms $\{a\} \sqsubseteq C$ and $\{a\} \sqsubseteq \exists r.\{b\}$, respectively. Those methods allowing for role chain and role inclusion [55], [68] incorporate additionally RBox constructs. In the case that the ABox is retained, assertional axioms are either embedded alongside with terminological axioms [55], [69], [96] or placed into the latent space after the model of TBox is constructed [70]. In contrast with other approaches, the method in [71] does not explicitly represent the geometric interpretations of individuals, roles and concepts for \mathcal{ELH} ontologies: it constructs a binary vector $\mu(d)$ of size $|\mathcal{N}_I| + |\mathcal{N}_C| + |\mathcal{N}_R| \cdot |\Delta^{\mathcal{I}}|$ associated with each domain element d such that $\mu(d)[a] = 1$ if $d = a^{\mathcal{I}}$ (for individuals), $\mu(d)[A] = 1$ if $d \in A^{\mathcal{I}}$ (for classes)

and $\mu(d)[r, e] = 1$ if $(d, e) \in r^{\mathcal{I}}$ (for relations) ($\mu(d)[a] = 0$, $\mu(d)[A] = 0$, $\mu(d)[r, e] = 0$ otherwise).

In order to show that learned embeddings construct a logical geometric model of a given ontology, many methods provide an explicit proof. In most cases, theoretical results state that if the optimization objective converges to a certain value during training and some other conditions are satisfied, the theory has a model [22], [55], [69], [96], and this model is constructed through the optimization process; in this sense, faithfulness (machine learning) holds. Sometimes, authors refer to this property as *soundness* [22], [55], [96]. For \mathcal{ELH} ontology embeddings [71] and axis-aligned cone embeddings [70] our definition of faithfulness is not applicable since in these works only interpretation domains and interpretation functions are introduced without describing optimization process which approximates geometric models. Definitions of strong and weak faithfulness discussed in these works are properties of an interpretation function \mathcal{I} .

As for works that do not provide theoretical proofs, EmEL++ [68] cannot faithfully capture the role hierarchy although the corresponding loss is present: the $r \sqsubseteq s$ loss is symmetric with respect to r and s , so $r \sqsubseteq s$ necessarily implies $s \sqsubseteq r$; the same holds for role chains $r_1 \circ r_2 \sqsubseteq s$ and $r_2 \circ r_1 \sqsubseteq s$. ELBE [54] has faithfulness property which relies

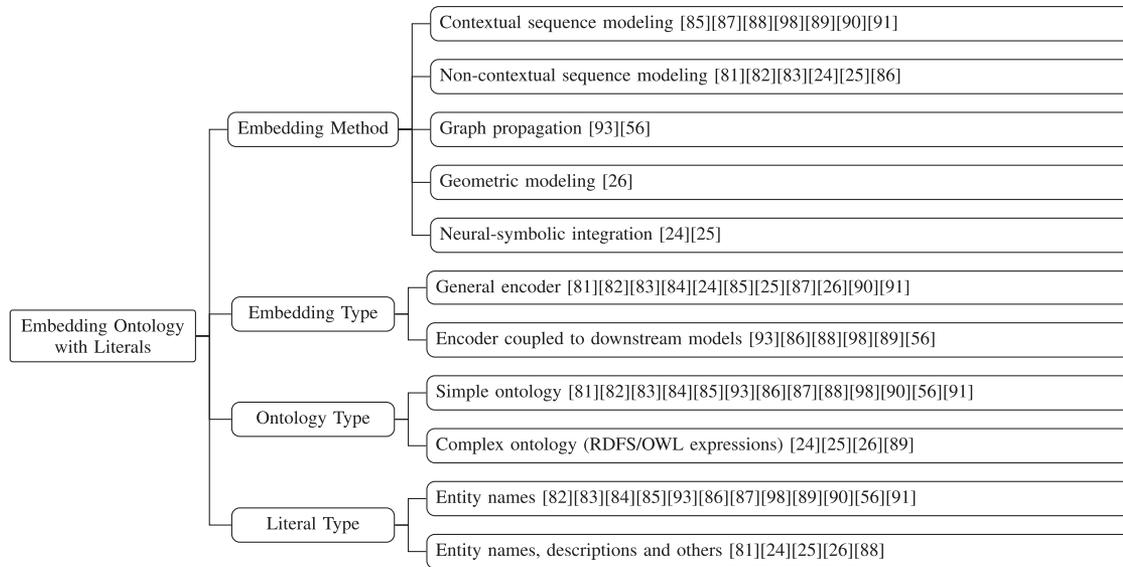


Fig. 6. Dimensions, their values and corresponding works of embedding ontology with literals.

on justifications discussed in [22] since each of its normal form losses is a multi-dimensional version of corresponding EL embeddings objectives; faithfulness (machine learning) in this particular case can be formulated as follows (using notations from the paper): if $margin$ is a vector with non-positive components and the total loss is equal to 0 then the model is constructed. CatE [72] does not generate models, yet it can be considered as faithful with respect to the lattice of ontology concepts in the sense that the transformation from the ontology to the lattice is total and injective, and the embedding preserves lattice structure.

Based on the analysis of these geometric models for complex ontology embedding, we have the following discussions:

- These models adopt highly interpretable balls and axis-aligned boxes as well as fuzzy sets and order embeddings, supporting constructs that are not covered by simple ontology embedding models [55], [68], [70], [71], [96], and contributing more fine-grained ontology embeddings [54], [69], [72].
- In terms of expressivity, these models are tailored to main DL constructs of either \mathcal{EL}^{++} [22], [54], [55], [68], [69], [71] or \mathcal{ALC} [70], [72], [96]. As one potential future direction, we can mention targeting the full expressivity of \mathcal{EL}^{++} , motivated by real-world applications (e.g., many ontologies in biomedical domain) and stronger reasoning capabilities.
- Applicability to real-world scenarios remains an open question for some embedding methods [71], [70], requiring more evaluation on real-world ontologies and tasks.

D. Embedding Ontology With Literals

We analyse those works for embedding ontologies with literals from four dimensions — embedding method, embedding type, ontology type and literal type. Their potential values as well as the corresponding works are shown in Fig. 6. For embedding type, we divide the methods into two kinds: general encoders

which are separately trained and applicable to different downstream tasks, and coupled encoders which are jointly trained with some other attached models often for feature learning and cannot be applied without these models. For literal type, we classify the works into those only use entity name information, and those use entity name information and other literals such as long descriptions.

Most methods for embedding ontologies with literals adopt the solution of sequential modeling demonstrated in Fig. 2. The biggest challenge of this solution lies in its first step — extracting literal-injected sequences from the ontology. Some methods such as ERSOM [81], DeepAlignment [83], \mathcal{N} -ball Embeddings [84], SapBERT [87] and HiT [91] directly use the literals of an entity like names and descriptions as its sequences to learn its embedding. Such *literal-alone sequences* miss the formal semantics of the entity, and therefore, some other methods such as BERTSubs [89], SORBERT [90] and [85] extract *context-augmented literal sequences* by first exploring the serialization of an entity’s contexts such as the axioms this entity is involved and its neighbourhood in a graph transformed from the ontology, and then textualising the entity sequences by replacing (a part) of the entities by their names. For deeper integration of the literals and the formal semantics, some more complex strategies, such as merging corpora of different kinds of sequences, and concatenating a literal-alone sequence and a context-augmented literal sequence of an entity for a *hybrid literal sequence*, have also been proposed in OPA2Vec [24] and OWL2Vec* [25].

In the second step of this sequence learning solution, many methods directly train an encoder from the extracted sequences by unsupervised learning. ERSOM [81] trains a stacked auto-encoder which is a neural network with several hidden layers; DeepAlignment [83], OPA2Vec [24] and OWL2Vec* [25] directly train a Word2Vec model; [85] trains a biomedical variant of BERT. There are also some methods trying to utilize some external tasks and data for training. SORBERT [90] trains a sentence transformer with a siamese network architecture by minimizing the distance between matched concepts from two ontologies; SapBERT [87] trains a biomedical BERT by

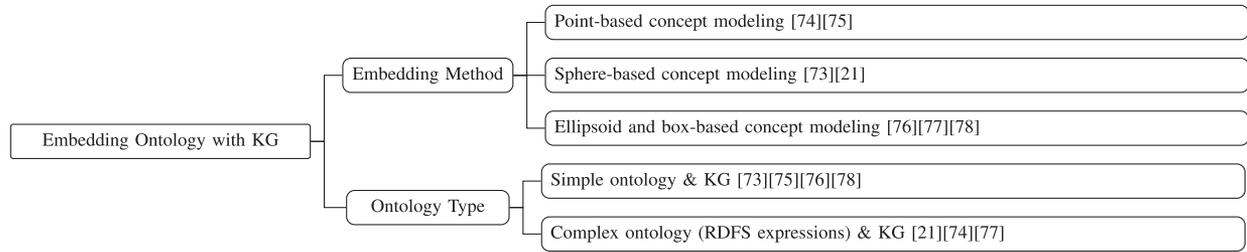


Fig. 7. Dimensions, their values and corresponding works of embedding ontology with KG.

minimizing the distance between a mention from text and its matched entity in an ontology; HiT [91] re-trains variants of BERT by retaining the subsumption relationships of an ontology in a Poincaré ball.

The encoders by the above two kinds of methods are usually quite general, i.e., they are applicable to different tasks and/or different other ontologies beyond the ones for training, and their embeddings can be fed to different other models. On the other hand, some encoders for ontology embedding are trained in conjunction with some downstream models by specific tasks. BERTSubs [89] fine-tunes a BERT model and an attached classifier with concept subsumptions; OntoProtein [88] trains a BERT model in conjunction with learning the embeddings of proteins and the Gene Ontology. Such learned ontology embeddings usually have limited generality, and can only be applied with another jointly trained model.

Besides sequence learning, OntoZSL [26] embeds an RDFS ontology with entity names and descriptions by a geometric modeling method which extends TransE with additional losses between entity representations and literal representations. MEDTO [93] and [56] both use Graph Neural Networks to learn concept features via propagation over concept hierarchies for minimizing distances between matched concepts. OPA2Vec [24] and OWL2Vec* [25] also adopt neural-symbolic integration by employing OWL reasoners for inferring hidden axioms for augmenting the sequence extraction.

With the above analysis on the current works, we have the following discussion for ontology embedding with literals:

- Sequence modeling is the most widely used general solution. Among these works, using contextual word embedding often leads to better performance in downstream tasks, but it requires more future efforts to train general encoders.
- Geometric modeling has been explored only in [26]. However, it can train embeddings with higher interpretability, through which we can understand how knowledge are inferred and learn the impact of semantics of literals.
- Current works focus on textual literals such as short phrases of names and long text of descriptions. Literals of other data types such as numbers are simply regarded as plain text. More future efforts are worthy while to deal with literals of different types especially images for more comprehensive multi-modal ontology embedding.

E. Embedding Ontology With KG

We analyze those works for embedding ontologies with KGs from two dimensions — embedding method and ontology type.

Note we consider those works that use the KG for supporting ontology embedding or jointly embed the KG and the ontology, but ignore those works that only use ontologies for supporting KG embedding (e.g., [99], [100]). See [32] for a survey of the latter. As all the works under consideration use geometric modeling, we make a more fine-grained categorization according to how concepts are modeled. Their potential values and corresponding works are shown in Fig. 7.

Some works consider joint embedding of hierarchical concepts and a KG with relational facts [73], [75], [76], [78]. TransC [73] represents each concept as a sphere and each instance as a point in the euclidean space. The concept membership is then modeled by point inclusion in sphere, and the concept subsumption is modeled by sphere inclusion. Their losses are jointly minimized together with the translation loss of TransE that models the facts. For higher expressivity, some more complex geometric objects are used. TransEllipsoid [76] and EIKE [78] both model each concept as a high dimensional ellipsoid, and TransCuboid [76] models each concept as a high dimensional box. The ellipsoid and box are both modeled by one vector for the center and another vector for the boundary (i.e., offset). TransEllipsoid, TransCuboid and EIKE use similar losses as TransC for training. OntoEA [75] jointly embeds two KGs and their ontologies. It uses a point to represent each concept, and models not only concept subsumption, instance membership and relational facts, but also instance matching across KGs and concept disjointness.

The other works consider embedding of RDFS ontologies with KGs [21], [74], [77]. JOIE [74] transforms the RDFS ontology into an ontology view graph, with each concept modeled as a node, and each relation's domain concepts and range concepts connected by this relation. In training, triples of the ontology view graph are equally modeled as normal KG triples by a translation loss, and the instance membership is modeled by a mapping from the instance to the concept. Concept2Box [77] extends JOIE by modeling each concept as a box and by replacing the translation loss with a binary cross entropy loss. EmbedS [21] models each concept as a sphere and each relation by two spheres — one for its domain and the other for its range, and uses distance-based losses. However, EmbedS has not been evaluated.

With the above analysis, we have the following observations and perspectives on embedding ontology with KG:

- Although the technical solutions of sequence modeling and graph propagation have not been exploited in the current works, through transforming the ontology and the KG into one graph or directly into sequences, they can be applied.

TABLE II
TIME AND SPACE COMPLEXITY OF ONTOLOGY EMBEDDING METHODS

Method	Space Complexity	Time Complexity
Simple Ontology		
Order Embeddings [57]	$\mathcal{O}(d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Poincaré Embeddings [18]	$\mathcal{O}(d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Hyperbolic Entailment Cones [62]	$\mathcal{O}(d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Box Lattices [19]	$\mathcal{O}(2d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Density Order Embeddings [58]	$\mathcal{O}(2d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Smooth Boxes [59]	$\mathcal{O}(2d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Joint Hierarchies with Boxes [60]	$\mathcal{O}(4d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
Gumbel Box Embeddings [61]	$\mathcal{O}(2d \cdot \mathbf{N}_C)$	$\mathcal{O}(d)$
HBE [65]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_R))$	$\mathcal{O}(d^2)$
HYPON [64]	$\mathcal{O}(d \cdot \mathbf{N}_I)$	$\mathcal{O}(d^2)$
HyperExpan[92]	$\mathcal{O}(d \cdot \mathbf{N}_C)$	$\mathcal{O}(d^2)$
Complex Ontology		
ELEmbeddings [22]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R) + \mathbf{N}_C + \mathbf{N}_I)$	$\mathcal{O}(d)$
EmEL++ [68]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R) + \mathbf{N}_C + \mathbf{N}_I)$	$\mathcal{O}(d)$
ELBE [54]	$\mathcal{O}(d \cdot (2 \mathbf{N}_C + 2 \mathbf{N}_I + \mathbf{N}_R))$	$\mathcal{O}(d)$
BoxEL [69]	$\mathcal{O}(d \cdot (2 \mathbf{N}_C + \mathbf{N}_I + 2 \mathbf{N}_R))$	$\mathcal{O}(d)$
Box ² EL [55]	$\mathcal{O}(d \cdot (3 \mathbf{N}_C + 2 \mathbf{N}_I + 4 \mathbf{N}_R))$	$\mathcal{O}(d)$
Convex/Non-convex Regions [71]	$\mathcal{O}(\Delta^{\mathcal{I}} \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R \cdot \Delta^{\mathcal{I}}))$	$\mathcal{O}(\Delta^{\mathcal{I}} ^5)$
Al-Cones [70]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I) + d^2 \cdot \mathbf{N}_R)$	$\mathcal{O}(d^2)$
FALCON [96]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_{I_e} + \mathbf{N}_R))$	$\mathcal{O}(d^2 \cdot \Delta^{\mathcal{I}})$
CatE [72]	$\mathcal{O}(d \cdot op)$	$\mathcal{O}(d)$
Ontology with KGs		
TransC [73]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R) + \mathbf{N}_C)$	$\mathcal{O}(d)$
EmbedS [21]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + 2 \mathbf{N}_R) + \mathbf{N}_C + \mathbf{N}_R)$	$\mathcal{O}(d)$
JOIE [74]	$\mathcal{O}(d_1 \cdot (\mathbf{N}_C + \mathbf{N}_R) + d_2 \cdot (\mathbf{N}_I + \mathbf{N}_R))$	$\mathcal{O}(d_1^2 + d_1 d_2)$
OntoEA [75]	$\mathcal{O}(d \cdot (\mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R))$	$\mathcal{O}(d^2)$
TransEllipsoid/TransCuboid [76]	$\mathcal{O}(d \cdot (2 \mathbf{N}_C + \mathbf{N}_I + \mathbf{N}_R))$	$\mathcal{O}(d)$
Concept2Box [77]	$\mathcal{O}(d \cdot (2 \mathbf{N}_C + 3 \mathbf{N}_R + \mathbf{N}_I))$	$\mathcal{O}(d \cdot d_{BERT})$
EIKE [78]	$\mathcal{O}(d \cdot (2 \mathbf{N}_C + \mathbf{N}_I) + \mathbf{N}_R)$	$\mathcal{O}(d^2)$
Ontology with Literals		
ERSOM [81]	$\mathcal{O}(V \cdot (\mathbf{N}_C + \mathbf{N}_R + \mathbf{N}_I))$	$\mathcal{O}(L \cdot d^2)$
DeepAlignment [83]	$\mathcal{O}(d \cdot V)$	$\mathcal{O}(d)$
Category Trees [84]	$\mathcal{O}(\log_b(V))$	$\mathcal{O}(\log_b(V))$
MEDTO [93]	$\mathcal{O}(d \cdot \mathbf{N}_C)$	$\mathcal{O}(L \cdot d^2)$
Semantic/Structural Embeddings [56]	$\mathcal{O}(w \cdot l \cdot V + d \cdot V)$	$\mathcal{O}(w \cdot l \cdot V + c \cdot d \cdot \log_2(V) + d^2)$
OPA2Vec [24]	$\mathcal{O}(d \cdot V)$	$\mathcal{O}(c \cdot d \cdot \log_2(V))$
OWL2Vec* [25]	$\mathcal{O}(w \cdot l \cdot V + d \cdot V)$	$\mathcal{O}(w \cdot l \cdot V + c \cdot d \cdot \log_2(V))$

- The current works adopt typical KG completion benchmarks for evaluation. More real-world benchmarks and complex tasks are required to full evaluation. Meanwhile, some complex ontology embedding methods such as Box²EL [55] are also applicable, but have not been evaluated.

F. Complexity of Ontology Embedding Methods

We analyze the space and time complexity of ontology embedding methods in Table II. The notation \mathbf{N}_C , \mathbf{N}_R , \mathbf{N}_I represents the set of concepts, role and individual names within an ontology, respectively, and d is the embedding dimension. Space complexity describes the number of memory units to store embedding vectors. Most of the methods increase linearly with d because ontology entities are represented as vectors in \mathbb{R}^d . However, some methods, like the one in [70] exhibit a quadratic space complexity, $\mathcal{O}(d^2)$, because role entities are stored as matrices in $\mathbb{R}^{d \times d}$. Particular cases include [71], where $d = |\mathbf{N}_C| + |\mathbf{N}_I| + |\mathbf{N}_R| \cdot |\Delta^{\mathcal{I}}|$ and each element in the domain $\Delta^{\mathcal{I}}$ is mapped to a binary vector v of length d and [72],

where complexity scales linearly with the number of operators op in ontology axioms. Operators are symbols $\sqcap, \sqcup, \neg, \exists, \forall$ found in ontology axioms. Since \mathcal{ALC} allows arbitrarily long axioms, we assume that $\mathcal{O}(d \cdot op) \gg \mathcal{O}(d \cdot (|\mathbf{N}_C| + |\mathbf{N}_R| + |\mathbf{N}_I|))$ in most cases.

Time complexity $\mathcal{O}(d)$ involves linear operations such as scalar multiplication or element-wise vector operations. In contrast, methods with complexity $\mathcal{O}(d^2)$ involve quadratic operations which usually take the form of matrix-vector multiplications. A particular case is [96], where the complexity is $\mathcal{O}(d^2 \cdot |\Delta^{\mathcal{I}}|)$ because concept descriptions involving existential or universal restrictions involve aggregation operations over all elements in the domain $\Delta^{\mathcal{I}}$.

Ontology embeddings with literals are particularly different as the set of entities include not only concepts, individuals and roles, but also a potentially large vocabulary from descriptions, labels or external documents. Therefore, the complexity analysis incorporates the vocabulary size $|V|$ and we assume that $|V| \geq |\mathbf{N}_C| + |\mathbf{N}_I| + |\mathbf{N}_R|$. Methods in [24] and [25] incorporate Word2Vec, whose space and time complexity are $\mathcal{O}(d \cdot |V|)$ and $\mathcal{O}(c \cdot d \cdot \log_2(|V|))$, respectively, where c represents the context

size in Word2Vec. The complexity of methods that implement random walks [25], [56] include parameters such as number of walks w and walk length l . We use the parameter L to represent the number of layers in a neural networks used in methods such as [81] and [93]. Finally, we do not include methods that involve training/fine-tuning language models because their large number of parameters and training time can obscure the complexity analysis.

IV. ONTOLOGY EMBEDDINGS FOR KNOWLEDGE ENGINEERING AND MACHINE LEARNING

A. Knowledge Engineering

1) *Ontology Matching*: Given two ontologies \mathcal{O}_1 and \mathcal{O}_2 , ontology matching (OM) is to find out entity mappings in form of (e_1, e_2) , where e_1 and e_2 are from \mathcal{O}_1 and \mathcal{O}_2 , respectively, with an equivalence or subsumption relationship [101]. A good OM system is expected to have a high Precision for the discovered mappings and a high Recall towards the ground truth mappings. Sometimes, some entities in one ontology are given, and for each of them, an OM system is expected to rank the entities in the other ontology such that the truly matched entity is ranked in the first position. In this situation, ranking-based metrics like MRR (Mean Reciprocal Rank) and Hits@K (K=1, 5, 10,...) are often used for evaluation.

Traditional systems mostly use some of the following three techniques: lexical matching, graph structure matching and logical reasoning. However, they are limited in several aspects such text understanding and fusion of different semantics. Ontology embeddings provide a promising solution to address these limitations, and thus have been applied for OM by several recent studies. Meanwhile, OM benchmarks that are specifically developed for evaluating machine learning-based OM systems such as Bio-ML [102] provide good contexts for evaluating ontology embeddings.

Most embedding-based OM methods consider literals due to their important information. The early methods ERSOM [81] and DeepAlignment [83] directly calculate the distance of two concept embeddings, while LogMap-ML [103] and SORBERT [90] further trains a supervised mapping classifier that uses concept embeddings as input. Both solutions exploit the embedded semantics for discovering mappings, but the improvement over the traditional systems is still limited as the embeddings are general with no specification to OM. For better performance, most recent OM methods including MEDTO [86], [93], BERTMap [98], BERTSubs [89] and [56] jointly learn task specific embeddings of an ontology with a model for matching. For example, BERTMap [98] fine-tunes a PLM for encoding concepts with their names, using synonyms from the two ontologies and the given mappings in option, while BERTSubs [89] fine-tunes a PLM for encoding concepts with their contexts for predicting concept subsumption mappings.

2) *Ontology Reasoning*: The embeddings of an ontology can be used to infer its missing knowledge, among which different forms of concept subsumptions such as $C \sqsubseteq D$, $C \sqcap D \sqsubseteq E$, $C \sqsubseteq \exists r.D$ and $\exists r.D \sqsubseteq C$, concept memberships, property domains and ranges are commonly considered by many current studies (e.g., [18], [25], [55], [68], [69], [74] and [89]). In evaluation for concept subsumption inference, the sub-concept

is often given, and a set of candidate concepts are ranked according to the score of being the super-concept, where MRR and Hits@K are often adopted for performance measurement. Note the selection of candidate concepts can be quite flexible, depending on the benchmarking requirement. For example, they can be all the named or complex concepts that exist in the ontology, a particularly selected subset of them, or some particularly constructed complex concepts. The evaluation for concept membership inference is similar with an instance given and a set of candidate concepts ranked.

Meanwhile, there are two settings for inference:

- *Prediction*: A small part of the axioms are splitted out from all the declared axioms of the ontology for testing, and the remaining declared axioms are used for training. The models are expected to capture more generalizable patterns for achieving better prediction performance.
- *Approximate Deductive Inference*: The declared axioms are used for training, while the entailed axioms are used for testing. This setting is often used for measuring whether the ontology embeddings have retained all the formal semantics.

Embeddings learned by different solutions are applied for ontology reasoning using different paradigms. (i) For embeddings by geometric modeling, axioms can usually be inferred by calculating the geometric relationship of vectors. For the embeddings that represent concepts by boxes, $C \sqsubseteq D$ can be inferred if the box of C is fully inside the box of D . Otherwise, a score can also be calculated according to the relative volume of their overlap. (ii) Embeddings by sequence learning and graph propagation can be regarded as pre-trained machine learning features and can be fed into another separately trained machine learning models like binary classifiers for concept subsumption prediction (e.g., [25], [85]) and unsupervised clusters for concept clustering (e.g., [104]).

3) *Discussion*: We have the following observations for ontology embedding for knowledge engineering. (i) Embeddings by geometric modeling support interpretable inference, but often perform worse than embeddings by sequence modeling with literals incorporated. It is challenging but promising to incorporate literals in geometric modeling. (ii) Most evaluation assumes a part of the knowledge to infer (e.g., the sub-concept of a concept subsumption axiom) are given, but such settings still require much human support in real-life scenarios. We need more benchmarks and metrics for supporting end-to-end evaluation. (iii) The application of ontology embeddings for knowledge engineering mostly lie in OM and inferring missing knowledge within ontology. Other tasks such as entity resolution, query answering, knowledge retrieval and ontology learning from text can be explored.

B. Knowledge Augmented Machine Learning

Ontologies are able to represent information of machine learning tasks, datasets and algorithms, and thus ontology embedding can be a medium to inject domain knowledge into machine learning training or prediction. One typical aspect for augmentation is dealing with the sample shortage problem [105], [106]. In this part, we introduce a case study of using ontology embeddings for zero-shot learning (ZSL) which typically refers to a machine

learning classification task⁷ with some or all of its testing labels unseen in training [27], [107]. The model is expected to have high accuracy on testing samples of both seen and unseen labels.

Ontology-aware ZSL requires to construct or re-use an ontology that models the relationships of seen and unseen classification labels, where each label is often represented as a concept in the ontology. For example, in animal image classification, such an ontology could represent animal taxonomies, visual characteristics, habitats, and so on. With the ontology and its embeddings, there are mainly two paradigms among the current studies:

- *Mapping-based*: In training, this paradigm learns a mapping function from the training samples to map the vector representation (e.g., image features) of the input to the ontology embedding of the output label. In prediction, the mapping function is applied to map the test input into an embedding, and the label (either seen or unseen) whose embedding is closest to this embedding is regarded as the output. It can also map the label's embedding to the input's vector, or map both to a common vector. One example work is [108] which uses EL Embedding [22] to embed an ontology of OWL EL for modeling labels of animals for zero-shot animal image classification.
- *Generative*: This paradigm generates samples for an unseen label according to its embedding in the ontology, by learning a conditional generative model like Generative Adversarial Network (GAN). Thus the ZSL problem is transformed into a normal supervised learning problem. A representative work is OntoZSL [26] which embeds literal-aware ontologies for zero-shot image classification, and zero-shot KG link prediction with unseen relations.

Applying ontology embedding for machine learning sample shortage is a promising solution of neural-symbolic integration, but there is still a shortage of successful systems under deployment. We think the limitation lies in the representation of more complex knowledge such as uncertain relationships (e.g., an image of horse may have a background of grass land, but not always), as well as the automatic construction of the ontology for a specific task. The corresponding solutions could be more flexible multi-modal ontologies with different kinds of literals and example instances, and more tools for knowledge integration and automatic ontology construction.

V. ONTOLOGY EMBEDDING DOMAIN APPLICATIONS

A. Life Sciences

In life sciences, the development of successful ontologies such as the Gene Ontology (GO) [9], SNOMED-CT [40] and the Human Phenotype Ontology [109] has motivated the development of methods that incorporate ontologies as a source of background knowledge. With the rise of machine learning, ontology embedding has become a common approach to leverage the ontology. In this part, we review works of two kinds of life science tasks which widely use ontology embeddings: (i) protein-protein

interaction, gene-disease association and protein/gene function prediction, and (ii) healthcare predictive analysis with Electronic Health Records (EHR).

Methods for these tasks are usually evaluated as (i) classification systems with metrics such as Precision and Recall, (ii) ranking systems with metrics like MRR and Hits@K, or (iii) predictive systems with metrics like Accuracy@K (K=5,20,...). In particular cases, methods are evaluated on problem-specific metrics. For example, protein function prediction uses F_{\max} which is obtained from Precision and Recall scores, and S_{\min} which computes the level of uncertainty and misinformation of the predictions.

The central idea of generating ontology embeddings is to capture an entity's latent relationships with other entities. For example, in protein-protein interaction, whose objective is to predict if two proteins interact, methods such as [24] and [88] link protein entities to their corresponding functions in the Gene Ontology, and thus the generated embeddings of protein entities encode information about their functions that can be utilised for interaction prediction. A similar idea is followed in the gene-disease association problem. For example, in [24], [119], [120] and [122], genes and diseases are linked to their corresponding phenotypes in a phenotype ontology, with the goal of capturing phenotypic-related information in the embeddings. Other strategies, such as [118], generate only phenotype embeddings and then compute an embedding for a gene (or disease) using their associated phenotype embeddings. In protein function prediction, embeddings for functions are obtained from GO, whereas embeddings for proteins are obtained from protein sequences, to which the use of PLMs is currently predominant [128], [129].

In healthcare, there are various EHR predictive analysis tasks such as mortality prediction, next-admission diagnosis prediction or hospital readmission prediction where hierarchical medical concepts are exploited: an ontology is compiled into a directed acyclic graph composed of its concepts and its embedding is combined with the textual information from EHRs. In most works including [110], [111], [112], [116], [113], [114], [115] and [121], the graph-based attention mechanism is applied to leverage the concept hierarchies; another approach is described in [63] and [121] where concept hierarchies are embedded through hyperbolic embedding. Embedded medical concepts can then be used for training a Recurrent Neural Network (RNN) for sequential diagnosis prediction or mortality prediction. [117] uses a dual RNN with co-attention and max pooling to fuse medical concept hierarchies of patient diagnoses and drugs for prescription recommendation. Several approaches use embeddings from multiple ontologies or multiple representations of a single ontology: [113] combines multi-relational ontologies via a graph attention network for multi-relational ontology embedding; [115] assigns multiple embeddings for non-leaf nodes (except for the root) of the ontology's directed acyclic graph.

The strategies that have been employed to utilise ontology embeddings for life science can be summarised as follows:

- *Similarity-based strategy*: Given the trained embeddings of a pair of entities, this strategy computes the pair's score by either directly calling a similarity function or using a neural network that is additionally trained for prediction. One

⁷ In machine learning classification, the output is often called class. To distinguish it with class in ontology, we call the output classification label or label in brief.

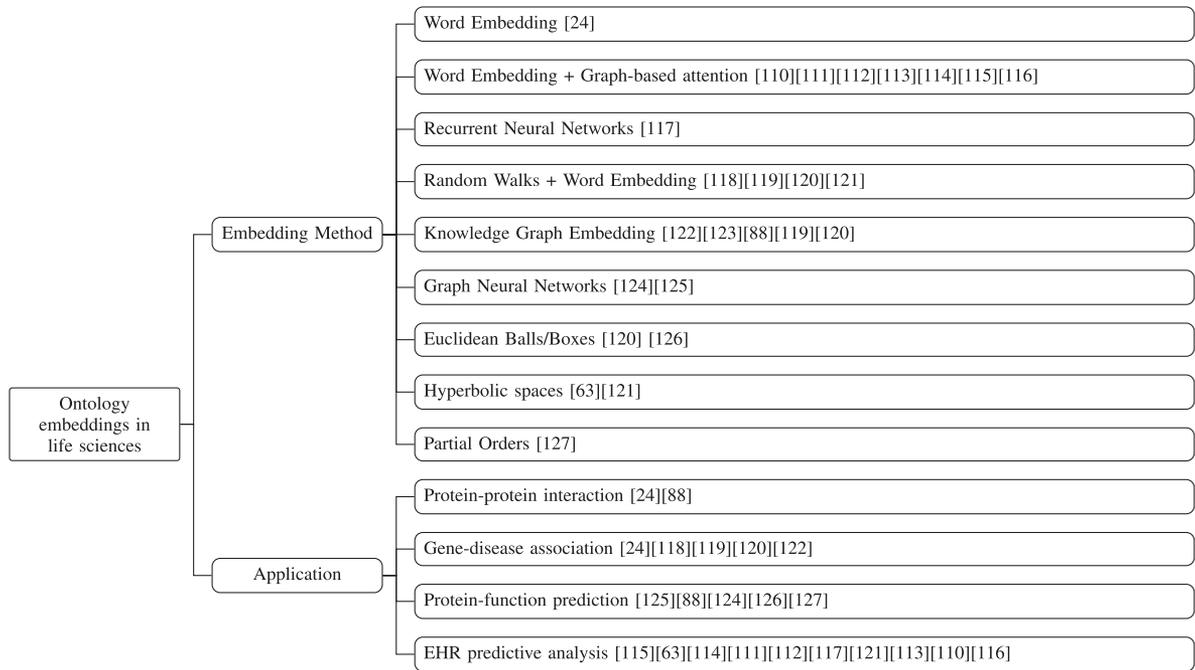


Fig. 8. Categorization of ontology embedding works for life sciences.

typical embedding method that is often applied in this strategy is OPA2Vec [24], which is to predict protein-protein interactions and gene-disease associations. More works that adopt this strategy include [118], [119] and [120]. Their embeddings are also based on the sequence learning model Word2Vec, trained with random walks of the graph extracted from the ontology.

- *Graph-based strategy*: This strategy is usually based on a graph created from one or many ontologies, applying techniques of KG embedding (KGE), neural networks, hyperbolic embedding, graph-based attention and so on. [88], [122] and [123] frame the problem as link prediction on the graph, and address it by KGE methods such as TransE [14] and DistMult [46]. OntoProtein [88] uses the graph to enhance the training of a Protein Language Model for better protein embeddings. PO2GO [127] embeds a directed acyclic graph from GO using a partial-order embedding method and also adopts an additional neural network for prediction. Works adopting GNNs either frame the problem as node classification [124] or use embeddings as partial input of a prediction module [125]. Graph-based attention models follow the original GRAM method [110] where the final embeddings of leaf concepts are convex combinations of their own embeddings and their ancestors' embeddings. More works include learning interaction of hierarchical embeddings of drug and diagnoses concepts by a dual RNN co-attention model [117] and utilizing the Poincaré ball model [63], [121].
- *Model-theoretic strategy*: Since OWL ontologies are formal semantics rooted in Description Logics, this strategy aims to utilize embeddings that are generated for theories. In this case, the prediction problem is framed as the inference of axioms with the embeddings. Methods such as [27]

for protein-function prediction and variations of [120] for gene-disease associations employ a model-theoretic embedding method ELEmbeddings [22]. [126] further generates multiple embedding models for approximate semantic entailment.

B. Other Applications

Although ontology utilization in other domains is not as prevalent as in life science applications, partially due to the absence of well-curated or widely used domain ontologies, there are some examples of ontology exploitation with embedding, including ontology-aware classifiers for identifying research topics in scholarly articles [130], enhancing intelligent transportation systems [131], sentiment analysis [132], event detection [133] and company cointegration prediction [134]. In these works, ontologies adopted are developed from scratch. Ontology embedding strategies in these selected studies are mostly quite simple or straightforward by applying some word or KG embedding methods, including Word2Vec in [130] and [131], XLNet for aspect-based sentiment extraction in [132], IterE [99] and Node2Vec [45] applied to ontology with KG in [133] and [134]. The work [135] discusses the task-dependent and task-independent evaluation of Poincaré disk embeddings for the GeoNames ontology [136].

VI. MOWL: A MACHINE LEARNING LIBRARY WITH ONTOLOGY EMBEDDING METHODS

Many ontology embedding works release the implementations of their methods or applications. However, researchers often face compatibility issues between different implementations, spend considerable time adapting code from various sources, and struggle to ensure fair comparisons between methods due to

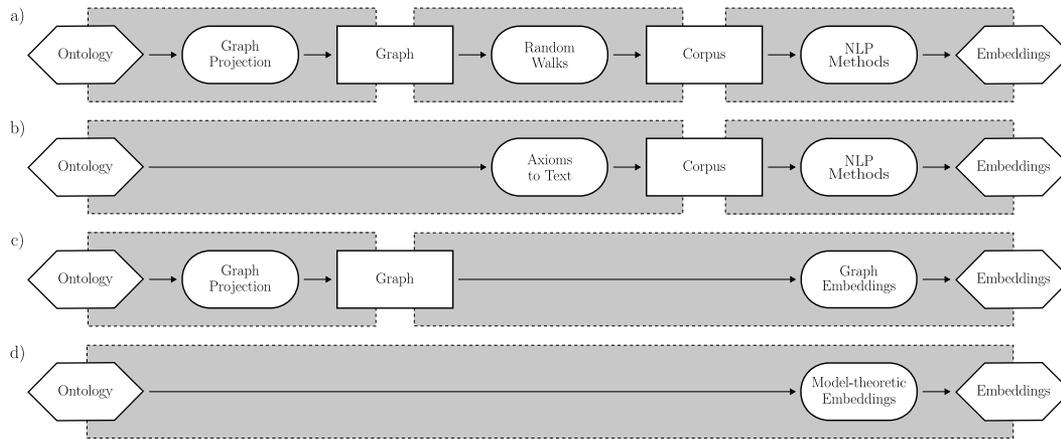


Fig. 9. High-level representation of four mOWL’s workflows for generating ontology embeddings, following the design patterns of neural-symbolic integration in [138]. In this representation, a grey box is a neural-symbolic design pattern that consume data, symbols or models and produce data symbols or models. Data and symbols are depicted with square boxes, models with hexahedrons and processes (such as training) with round boxes. Workflows a) and b) represent methods that transform the input ontology into sequences, and then learns the embeddings using NLP methods. These workflows are suitable for many methods of the sequence modeling solution. Workflow c) represents methods that transform the ontology into a graph and use graph embedding methods. This workflow is suitable for the graph propagation solution. Workflow d) directly uses an ontology’s axioms to learn the embeddings, which are suitable methods of geometric modeling, especially model-theoretic methods.

differences in implementation. Furthermore, there is a shortage of easy-to-use softwares that have implemented multiple ontology embedding methods and can support the implementation of new methods. mOWL⁸ [37], a library that provides functions to manipulate ontologies and use ontologies with machine learning, aims to bridge this gap. mOWL significantly reduces the engineering overhead for researchers by providing a unified API that standardizes common operations and allows fair comparison between methods. It enables researchers to focus on developing novel embedding approaches rather than dealing with infrastructure challenges. mOWL provides Python interfaces and can support the following functions: (i) ontology manipulation, where the OWL API [137] is accessed for ontology creation, manipulation and reasoning; (ii) ontology transformation, which enables extracting different graphs (such as concept hierarchies) and sequences from ontologies; (iii) implementation of ontology embedding methods that support ontologies of DL \mathcal{EL}^{++} and \mathcal{ALC} , covering required components of many methods described in Section III-A; (iv) common datasets and evaluation modules for axiom prediction and approximate deductive inference. mOWL includes several workflows following the modular design patterns for neural-symbolic systems in [138]: Fig. 9(a) and (b) encompass methods that transform ontology axioms and literals into sequences and use NLP methods to generate embeddings; Fig. 9(c) groups the methods that construct graphs from ontology axioms and leverage graph propagation methods to generate embeddings; Fig. 9(d) groups some model theoretic methods, especially the geometric methods targeting DL \mathcal{EL}^{++} . mOWL’s modular design makes it particularly suitable for both research and practical applications. Researchers can easily conduct comparative studies across different embedding methods, while practitioners can integrate mOWL into their workflows

To demonstrate mOWL, we use two different tasks and two benchmarks for each task. Both tasks adopt the prediction setting as described in Section IV-A2, which is to predict new axioms

from the existing ontology background knowledge. The first task *subsumption prediction* is to predict axioms are of the form $C \sqsubseteq D$ where C, D are concept names. The two benchmarks used for subsumption prediction are constructed from GO and the Food Ontology [25] and included in mOWL. The second task *protein–protein interaction prediction* (PPI) is to determine whether two proteins interact or not based on their biological functions in the Gene Ontology. This task is formulated as prediction of axioms $p_i \sqsubseteq \exists \text{interacts_with}.p_j$, where p_i, p_j are instances of proteins. We tested PPIs for yeast and human organisms. Both the tasks are framed as ranking problems, with metrics of Mean Rank (MR), Mean Reciprocal Rank (MRR), Hits@k and AUC of ROC curve. In Table III, we showcase eight ontology embedding methods including methods that involve literals (OPA2Vec, OPA2Vec-NN, OWL2Vec*), methods that rely on KGE (OWL2Vec*-TransE), geometric methods that target DL \mathcal{EL}^{++} (ELEmbeddings, BoxEL, Box²EL) and methods that target DL \mathcal{ALC} (CatE). Their implementations in mOWL leverages a standard interface which not only eases the implementation of ontology embedding methods, but also enables their manipulation, analysis and extension.

VII. CHALLENGES AND FUTURE DIRECTIONS

Although many studies have been done, ontology embedding is still a relatively new direction, and there are several challenges that prevent ontology embedding from having more real-world applications. We believe more future works are required in at least the following aspects.

Efficient Geometric Modeling for Complex Ontologies. On the one hand, quite a few geometric modeling methods in the euclidean space have been proposed, but they are mostly limited to the main features (constructs) of DL \mathcal{EL}^{++} and \mathcal{ALC} . Many other features like the at-least and at-most restrictions have not been explored, not to mention modeling the complete formal semantics of an arbitrary OWL ontology. Thus, we need to extend these methods to support more features that are used in

⁸ <https://github.com/bio-ontology-research-group/mowl>

TABLE III
EVALUATION OF ONTOLOGY EMBEDDING METHODS IN SUBSUMPTION PREDICTION AND PROTEIN-PROTEIN INTERACION PREDICTION, WITH THE DATASET AND MODELS IMPLEMENTED IN MOWL

Method	MR	MRR	H@3	H@10	H@100	AUC
Prediction of axioms $C \sqsubseteq D$ in GO						
OPA2Vec	440	0.112	0.112	0.207	0.571	0.990
OPA2Vec-NN	137	<u>0.167</u>	<u>0.177</u>	<u>0.353</u>	<u>0.780</u>	0.997
OWL2Vec*-Sim	<u>141</u>	0.220	0.245	0.428	0.811	0.997
OWL2Vec*-TransE	7446	0.036	0.037	0.065	0.161	0.832
ELEmbeddings	3279	0.041	0.045	0.111	0.330	0.926
BoxEL	6981	0.007	0.003	0.020	0.076	0.842
Box ² EL	3940	0.031	0.027	0.088	0.305	0.911
CatE	4548	0.069	0.081	0.216	0.433	0.897
Prediction of axioms $C \sqsubseteq D$ in FoodOn						
OPA2Vec	2094	0.081	0.082	0.136	0.349	0.926
OPA2Vec-NN	284	<u>0.112</u>	0.109	<u>0.253</u>	<u>0.645</u>	0.990
OWL2Vec*-Sim	433	0.233	0.267	0.442	0.731	<u>0.985</u>
OWL2Vec*-TransE	7909	0.048	0.052	0.078	0.145	0.719
ELEmbeddings	3088	0.081	0.102	0.168	0.293	0.891
BoxEL	3790	0.028	0.035	0.078	0.192	0.866
Box ² EL	4655	0.039	0.051	0.118	0.221	0.835
CatE	4542	0.084	<u>0.146</u>	0.197	0.297	0.839
Prediction of PPI (yeast) axioms of the form $p_i \sqsubseteq \exists \text{interacts_with} . p_j$						
OPA2Vec	396	0.061	0.051	0.128	0.543	0.935
OPA2Vec-NN	172	0.144	0.147	0.326	0.777	0.971
OWL2Vec*-Sim	197	0.149	0.154	0.301	0.730	0.967
OWL2Vec*-TransE	219	0.190	0.203	<u>0.402</u>	<u>0.793</u>	0.964
ELEmbeddings	289	0.101	0.094	0.252	0.730	0.952
BoxEL	231	0.037	0.021	0.073	0.551	0.962
Box ² EL	188	<u>0.167</u>	<u>0.190</u>	0.435	0.805	<u>0.969</u>
CatE	259	0.043	<u>0.025</u>	0.093	0.563	<u>0.957</u>
Prediction of PPI (human) axioms of the form $p_i \sqsubseteq \exists \text{interacts_with} . p_j$						
OPA2Vec	678	0.080	0.071	0.177	0.594	0.961
OPA2Vec-NN	390	0.136	0.139	0.285	<u>0.692</u>	0.978
OWL2Vec*-Sim	568	0.131	0.140	0.283	0.639	0.967
OWL2Vec*-TransE	477	0.173	0.187	0.357	0.717	0.972
ELEmbeddings	812	0.081	0.075	0.175	0.573	0.953
BoxEL	<u>411</u>	0.038	0.021	0.079	0.564	<u>0.976</u>
Box ² EL	564	<u>0.163</u>	<u>0.175</u>	<u>0.336</u>	0.683	0.967
CatE	492	0.059	0.043	0.136	0.629	0.972

We report filtered metrics that exclude axioms in the training set.

real-world ontologies, with faithfulness in option. On the other hand, the main contents of real-world ontologies are usually the hierarchical concepts, but their modeling in the euclidean space (e.g., by high dimensional boxes) is much less efficient (i.e., requires many more parameters to learn) than their modeling in some hyperbolic spaces such as Poincaré ball in which the distance of points increases exponentially as they get closer to the boundary [18]. There have been several studies to explore geometric modeling in hyperbolic spaces for ontology embedding, such as [18], [62], [65] and [92], but they mostly only embed concept hierarchies. Geometric modeling in hyperbolic spaces for some other ontology features deserves higher attention.

Utilising and Supporting Large Language Models (LLMs): LLMs like the GPT series and the Llama series have shown great success in understanding not only natural language but also images and (semi-)structured data [139], [140]. A promising idea is to embed ontologies, especially those with literals, using LLMs. There have been some ontology embedding methods that use encoder-based language models, such as OWL2Vec* [25] and OPA2Vec [24] which use Word2Vec, and HiT [91] which uses BERT-like Transformer-based encoders, but the

generative LLMs are quite different as they adopt some decoder or encoder-decoder architectures. Therefore, novel solutions, such as further pre-training and/or instruction tuning, need to be explored to apply them for ontology embedding. Meanwhile, LLMs also suffer from several problems including hallucination, black-box and shortage of domain knowledge. Integrating KGs as well as other (semi-)structured data is widely regarded as a promising solution [49], [141]. How can we use ontology embedding to integrate ontologies with LLMs? How to support Retrieval Augmented Generation (RAG) [142] with ontology embedding for incorporating domain knowledge and reasoning? Both questions are worthwhile for future ontology embedding exploration.

Application in Neural-symbolic Integration and Domains: Currently, ontology embedding is mostly applied to construct and curate ontologies themselves, and link prediction for domains like life sciences. As a popular semantic technique, ontology has a high potential for building neural-symbolic integration [143], utilizing ontology embedding for incorporating domain knowledge and reasoning capabilities in machine learning. Although some works have been proposed for ontology-based zero-shot and few-shot learning [105], we believe the research in

this direction is far from enough, with many topics like using ontologies for supporting meta learning (e.g., model selection) and augmenting explanation, not fully explored. Meanwhile, more domain applications, both in and out of life sciences, should be considered for exploring the potential of ontology embeddings. The current life science applications, which only consider link prediction of several simple relations and use formal semantics without literals, are quite simple. Link prediction with more complex target relations and data of different modalities can be explored using literal-aware ontology embedding methods. Other tasks, such as generation for drug discovery and protein design [144] and natural language inference for clinical trial [145], can also be explored with ontology embedding.

Benchmarking: This direction still lacks systematic benchmarking resources. Ontology construction and inference are straightforward and effective for evaluating different aspects of ontology embeddings, including how much formal and informal semantics the embeddings retain. But most of the current works limit the evaluation to concept subsumption inference and concept alignment. More complex tasks in ontology learning, either utilizing or not utilizing external resources, such as learning complex concept axioms, and inserting new concepts, have been rarely considered. Meanwhile, these aforementioned tasks for neural-symbolic integration and domain applications can be considered for benchmarking.

VIII. CONCLUSION

This is a comprehensive survey of ontology embedding which is to represent knowledge of ontologies in vector spaces with their semantics (partially) retained. It gives formal definitions and properties of ontology embedding, and summarizes three different technical solutions, i.e., geometric modeling, sequence modeling and graph propagation, and categorizes the studies according to not only the methods used but also the ontologies they aim at (including simple ontology, complex ontology in OWL or RDFS, ontology with literal and ontology with KG). Following the method part, the survey also gives a relatively complete analysis for the application of ontology embedding in knowledge engineering, life sciences and machine learning augmentation, and demonstrates a library mOWL developed by the co-authors that has implemented several typical ontology embedding methods and benchmarking resources. In the end, the survey discusses some potential future directions, including the interesting topics of integrating ontology embedding with LLMs.

ACKNOWLEDGMENT

All the authors participated the discussion, paper writing and proof reading. Olga contributed the review of complex ontology embedding (III-C), applications in life sciences (V-A) and other domains (V-B). Fernando contributed the complexity analysis of embedding methods (III-F), the review of applications in life sciences (V-A) and other domains (V-B), and the demonstration of mOWL (VI). Robert co-led the work and contributed the definitions and properties (II-B4). Yuan contributed the review of simple ontology embedding (III-B). Jiaoyan co-led the work and contributed to the other parts.

REFERENCES

- [1] N. Guarino, D. Oberle, and S. Staab, "What is an ontology?," in *Handbook on Ontologies*. Berlin, Germany: Springer, 2009, pp. 1–17.
- [2] Y. Raimond, T. Scott, S. Oliver, P. Sinclair, and M. Smethurst, "Use of Semantic Web technologies on the BBC web sites," in *Linking Enterprise Data*. Berlin, Germany: Springer, 2010, pp. 263–283.
- [3] J. E. Harrison, S. Weber, R. Jakob, and C. G. Chute, "ICD-11: An international classification of diseases for the twenty-first century," *BMC Med. Inform. Decis. Mak.*, vol. 21, pp. 1–10, 2021.
- [4] X. L. Dong, "Challenges and innovations in building a product knowledge graph," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2869–2869.
- [5] I. Horrocks, "Ontologies and the Semantic Web," *Commun. ACM*, vol. 51, no. 12, pp. 58–67, 2008.
- [6] J. Hendler, O. Lassila, and T. Berners-Lee, "The Semantic Web," *Sci. Amer.*, vol. 284, no. 5, pp. 34–43, 2001.
- [7] B. McBride, "The resource description framework (RDF) and its vocabulary description language RDFS," in *Handbook on Ontologies*. Berlin, Germany: Springer, 2004, pp. 51–65.
- [8] F. Baader, I. Horrocks, C. Lutz, and U. Sattler, *Introduction to Description Logic*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [9] G. O. Consortium, "The gene ontology resource: 20 years and still going strong," *Nucleic Acids Res.*, vol. 47, no. D1, pp. D330–D338, 2019.
- [10] D. M. Dooley et al., "FoodOn: A harmonized food ontology to increase global food traceability, quality control and data integration," *NPJ Sci. Food*, vol. 2, no. 1, 2018, Art. no. 23.
- [11] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, "DBpedia: A nucleus for a web of open data," in *Proc. 6th Int. Semantic Web, 2nd Asian Conf. Asian Semantic Web Conf.*, 2007, pp. 722–735.
- [12] F. J. García-Peñalvo et al., "Analysis of the OWL ontologies: A survey," *Sci. Res. Essays*, vol. 6, pp. 4318–4329, 2011.
- [13] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," 2013, *arXiv:1301.3781*.
- [14] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2787–2795.
- [15] P. Ristoski and H. Paulheim, "RDF2Vec: RDF graph embeddings for data mining," in *Proc. Int. Semantic Web Conf.*, 2016, pp. 498–514.
- [16] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Trans. Knowl. Data Eng.*, vol. 29, no. 12, pp. 2724–2743, Dec. 2017.
- [17] X. Chen, S. Jia, and Y. Xiang, "A review: Knowledge reasoning over knowledge graph," *Expert Syst. Appl.*, vol. 141, 2020, Art. no. 112948.
- [18] M. Nickel and D. Kiela, "Poincaré embeddings for learning hierarchical representations," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6341–6350.
- [19] L. Vilnis, X. Li, S. Murty, and A. McCallum, "Probabilistic embedding of knowledge graphs with box lattice measures," in *Proc. Conf. Assoc. Comput. Linguistics*, 2018, pp. 263–272.
- [20] G. A. Miller, "Wordnet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [21] G. I. Diaz et al., "Embeds: Scalable, ontology-aware graph embeddings," in *Proc. Int. Conf. Extending Database Technol.*, 2018, pp. 433–436.
- [22] M. Kulmanov, W. Liu-Wei, Y. Yan, and R. Hoehndorf, "EL embeddings: Geometric construction of models for the description logic EL," in *Proc. Int. Joint Conf. Artif. Intell.*, 2019, pp. 6103–6109.
- [23] D. Garg, S. Ikbal, S.K. Srivastava, H. Vishwakarma, H. Karanam, and L. V. Subramaniam, "Quantum embedding of knowledge for reasoning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 5594–5604.
- [24] F. Z. Smali, X. Gao, and R. Hoehndorf, "OPA2Vec: Combining formal and informal content of biomedical ontologies to improve similarity-based prediction," *Bioinformatics*, vol. 35, no. 12, pp. 2133–2140, 2019.
- [25] J. Chen, P. Hu, E. Jimenez-Ruiz, O. M. Holter, D. Antonyrajah, and I. Horrocks, "OWL2Vec*: Embedding of owl ontologies," *Mach. Learn.*, vol. 110, no. 7, pp. 1813–1845, 2021.
- [26] Y. Geng et al., "OntoZSL: Ontology-enhanced zero-shot learning," in *Proc. Web Conf.*, 2021, pp. 3325–3336.
- [27] M. Kulmanov and R. Hoehndorf, "DeepGOZero: Improving protein function prediction from sequence and zero-shot learning based on ontology axioms," *Bioinformatics*, vol. 38, no. Supplement_1, pp. i238–i245, 2022.
- [28] G. Stoilos, G. B. Stamou, V. Tzouvaras, J. Z. Pan, and I. Horrocks, "Fuzzy OWL: Uncertainty and the Semantic Web," in *Proc. Workshop OWL: Experiences Directions*, 2005, pp. 11–12.

- [29] A. Maedche and S. Staab, "Ontology learning for the Semantic Web," *IEEE Intell. Syst.*, vol. 16, no. 2, pp. 72–79, Mar./Apr. 2001.
- [30] J. Lehmann, "DI-learner: Learning concepts in description logics," *J. Mach. Learn. Res.*, vol. 10, pp. 2639–2642, 2009.
- [31] S. Badreddine, A. D. Garcez, L. Serafini, and M. Spranger, "Logic tensor networks," *Artif. Intell.*, vol. 303, 2022, Art. no. 103649.
- [32] W. Zhang, J. Chen, J. Li, Z. Xu, J. Z. Pan, and H. Chen, "Knowledge graph reasoning with logics and embeddings: Survey and perspective," 2022, *arXiv:2202.07412*.
- [33] M. Alam, F. van Harmelen, and M. Acosta, "Towards semantically enriched embeddings for knowledge graph completion," 2023, *arXiv:2308.00081*.
- [34] B. Xiong et al., "Geometric relational embeddings: A survey," 2023, *arXiv:2304.11949*.
- [35] M. Kulmanov, F. Z. Smaili, X. Gao, and R. Hoehndorf, "Semantic similarity and machine learning with ontologies," *Brief. Bioinf.*, vol. 22, no. 4, 2021, Art. no. bbaa 199.
- [36] F. Z. Smaili, X. Gao, and R. Hoehndorf, "Onto2Vec: Joint vector-based representation of biological entities and their ontology-based annotations," *Bioinformatics*, vol. 34, no. 13, pp. i52–i60, 2018.
- [37] F. Zhapa-Camacho, M. Kulmanov, and R. Hoehndorf, "mOWL: Python library for machine learning with biomedical ontologies," *Bioinformatics*, vol. 39 2022, Art. no. btac811.
- [38] A. Hogan et al., "Knowledge graphs," *ACM Comput. Surv.*, vol. 54, no. 4, pp. 1–37, 2021.
- [39] F. Baader, S. Brandt, and C. Lutz, "Pushing the EL envelope," in *Proc. Int. Joint Conf. Artif. Intell.*, 2005, pp. 364–369.
- [40] K. Donnelly et al., "SNOMED-CT: The advanced terminology and coding system for ehealth," *Stud. Health Technol. Inform.*, vol. 121, 2006, Art. no. 279.
- [41] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proc. Conf. Empir. Methods Natural Lang. Process.*, 2014, pp. 1532–1543.
- [42] A. Vaswani et al., "Attention is all you need," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.
- [43] J. D. M.-W. C. Kenton and L. K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Assoc. Comput. Linguistics*, 2019, pp. 4171–4186.
- [44] E. Asgari and M. R. Mofrad, "Continuous distributed representation of biological sequences for deep proteomics and genomics," *PLoS One*, vol. 10, no. 11, 2015, Art. no. e0141287.
- [45] A. Grover and J. Leskovec, "Node2Vec: Scalable feature learning for networks," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 855–864.
- [46] B. Yang, S. W.-T. Yih, X. He, J. Gao, and L. Deng, "Embedding entities and relations for learning and inference in knowledge bases," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [47] T. Dettmers, P. Minervini, P. Stenetorp, and S. Riedel, "Convolutional 2D knowledge graph embeddings," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 1811–1818.
- [48] G. A. Gesese, R. Biswas, M. Alam, and H. Sack, "A survey on knowledge graph embeddings with literals: Which model links better literal-ly?," *Semantic Web*, vol. 12, no. 4, pp. 617–647, 2021.
- [49] J. Z. Pan et al., "Large language models and knowledge graphs: Opportunities and challenges," *Trans. Graph Data Knowl.*, vol. 1, pp. 1–38, 2023.
- [50] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [51] *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge, U.K.: Cambridge Univ. Press, Aug. 2007.
- [52] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, 2020, Art. no. 18.
- [53] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," 2017, *arXiv: 1702.08608*.
- [54] X. Peng, Z. Tang, M. Kulmanov, K. Niu, and R. Hoehndorf, "Description logic EL embeddings with intersectional closure," 2022, *arXiv:2202.14018*.
- [55] M. Jackermeier, J. Chen, and I. Horrocks, "Dual box embeddings for the description logic EL," in *Proc. Web Conf.*, 2024, pp. 2250–2258.
- [56] Z. Hao, W. Mayer, J. Xia, G. Li, L. Qin, and Z. Feng, "Ontology alignment with semantic and structural embeddings," *J. Web Semantics*, vol. 78, 2023, Art. no. 100798.
- [57] I. Vendrov, R. Kiros, S. Fidler, and R. Urtasun, "Order-embeddings of images and language," in *Proc. 4th Int. Conf. Learn. Representations*, 2016.
- [58] B. Athiwaratkun and A. G. Wilson, "Hierarchical density order embeddings," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [59] X. Li, L. Vilnis, D. Zhang, M. Boratko, and A. McCallum, "Smoothing the geometry of probabilistic box embeddings," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [60] D. Patel and S. Sankar, "Representing joint hierarchies with box embeddings," in *Proc. Conf. Automated Knowl. Base Construction*, 2020, doi: [10.24432/CSKS37](https://doi.org/10.24432/CSKS37).
- [61] S. Dasgupta, M. Boratko, D. Zhang, L. Vilnis, X. Li, and A. McCallum, "Improving local identifiability in probabilistic box embeddings," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 182–192.
- [62] O. Ganea, G. Bécigneul, and T. Hofmann, "Hyperbolic entailment cones for learning hierarchical embeddings," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1646–1655.
- [63] Q. Lu et al., "Learning electronic health records through hyperbolic embedding of medical ontologies," in *Proc. ACM Int. Conf. Bioinf., Comput. Biol. Health Inform.*, 2019, pp. 338–346.
- [64] Z. Li and S. Wang, "HYPON: Embedding biomedical ontology with entity sets," in *Proc. ACM Int. Conf. Bioinf., Comput. Biol. Health Inform.*, 2021, pp. 1–7.
- [65] Z. Pan and P. Wang, "Hyperbolic hierarchy-aware knowledge graph embedding for link prediction," in *Proc. Findings Conf. Empir. Methods Natural Lang. Process.*, 2021, pp. 2941–2948.
- [66] V. Gutierrez Basulto and S. Schockaert, "From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules," in *Proc. Int. Conf. Princ. Knowl. Representation Reasoning*, 2018, pp. 379–388.
- [67] Ö. L. Özçep, M. Leemhuis, and D. Wolter, "Cone semantics for logics with negation," in *Proc. Int. Joint Conf. Artif. Intell.*, 2020, pp. 1820–1826.
- [68] S. Mondal, S. Bhatia, and R. Mutharaju, "EmEL : Embeddings for ϵ l description logic," in *Peroc. AAAI Spring Symp. Combining Mach. Learn. Knowl. Eng.*, 2021. [Online]. Available: <https://ceur-ws.org/>
- [69] B. Xiong, N. Potyka, T.-K. Tran, M. Nayyeri, and S. Staab, "Faithful embeddings for EL knowledge bases," in *Proc. Int. Semantic Web Conf.*, 2022, pp. 22–38.
- [70] Ö. L. Özçep, M. Leemhuis, and D. Wolter, "Embedding ontologies in the description logic ALC by axis-aligned cones," *J. Artif. Intell. Res.*, vol. 78, pp. 217–267, 2023.
- [71] V. Lacerda, A. Ozaki, and R. Guimarães, "Strong faithfulness for ϵ l ontology embeddings," 2023, *arXiv:2310.02198*.
- [72] F. Zhapa-Camacho and R. Hoehndorf, "Lattice-preserving ALC ontology embeddings," in *Proc. 18th Int. Conf. Neural-Symbolic Learn. Reasoning*, 2024, pp. 355–369.
- [73] X. Lv, L. Hou, J. Li, and Z. Liu, "Differentiating concepts and instances for knowledge graph embedding," in *Proc. Conf. Empir. Methods Natural Lang. Process.*, 2018, pp. 1971–1979.
- [74] J. Hao, M. Chen, W. Yu, Y. Sun, and W. Wang, "Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 1709–1719.
- [75] Y. Xiang, Z. Zhang, J. Chen, X. Chen, Z. Lin, and Y. Zheng, "OntoEA: Ontology-guided entity alignment via joint knowledge graph embedding," in *Proc. Findings Conf. Assoc. Comput. Linguistics*, 2021, pp. 1117–1128.
- [76] J. Yu, C. Zhang, Z. Hu, Y. Ji, D. Fu, and X. Wang, "Geometry-based anisotropy representation learning of concepts for knowledge graph embedding," *Appl. Intell.*, vol. 53, no. 17, pp. 19940–19961, 2023.
- [77] Z. Huang et al., "Concept2Box: Joint geometric embeddings for learning two-view knowledge graphs," in *Proc. Findings Conf. Assoc. Comput. Linguistics*, 2023, pp. 10105–10118.
- [78] K. Wang, G. Qi, J. Chen, and T. Wu, "Embedding ontologies via incorporating extensional and intensional knowledge," 2024, *arXiv:2402.01677*.
- [79] M. Alshahrani, M. A. Khan, O. Maddouri, A. R. Kinjo, N. Queralt-Rosinach, and R. Hoehndorf, "Neuro-symbolic representation learning on biological knowledge graphs," *Bioinformatics*, vol. 33, no. 17, pp. 2723–2730, 2017.
- [80] O. M. Holter, E. B. Myklebust, J. Chen, and E. Jiménez-Ruiz, "Embedding OWL ontologies with OWL2Vec," in *Proc. CEUR Workshop Proc.*, 2019, pp. 33–36.

- [81] C. Xiang, T. Jiang, B. Chang, and Z. Sui, "ERSOM: A structural ontology matching approach using automatically learned entity representation," in *Proc. Conf. Empir. Methods Natural Lang. Process.*, 2015, pp. 2419–2429.
- [82] V. Jayawardana, D. Lakmal, N. de Silva, A. S. Perera, K. Sugathadasa, and B. Ayesha, "Deriving a representative vector for ontology classes with instance word vector embeddings," in *Proc. 7th Int. Conf. Innov. Comput. Technol.*, 2017, pp. 79–84.
- [83] P. Kolyvakis, A. Kalousis, and D. Kiritsis, "DeepAlignment: Unsupervised ontology matching with refined word vectors," in *Proc. Conf. North Amer. Assoc. Comput. Linguistics*, 2018, pp. 787–798.
- [84] T. Dong et al., "Imposing category trees onto word-embeddings using a geometric construction," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [85] H. Liu, Y. Perl, and J. Geller, "Concept placement using BERT trained by transforming and summarizing biomedical ontology structure," *J. Biomed. Inform.*, vol. 112, 2020, Art. no. 103607.
- [86] V. Nguyen, H. Y. Yip, and O. Bodenreider, "Biomedical vocabulary alignment at scale in the UMLS metathesaurus," in *Proc. Web Conf.*, 2021, pp. 2672–2683.
- [87] F. Liu, E. Shareghi, Z. Meng, M. Basaldella, and N. Collier, "Self-alignment pretraining for biomedical entity representations," in *Proc. Conf. North Amer. Assoc. Comput. Linguistics*, 2021, pp. 4228–4238.
- [88] N. Zhang et al., "OntoProtein: Protein pretraining with gene ontology embedding," in *Proc. Int. Conf. Learn. Representations*, 2022.
- [89] J. Chen, Y. He, Y. Geng, E. Jiménez-Ruiz, H. Dong, and I. Horrocks, "Contextual semantic embeddings for ontology subsumption prediction," *World Wide Web*, vol. 26, no. 5, pp. 2569–2591, 2023.
- [90] F. Gosselin and A. Zouaq, "SORBET: A siamese network for ontology embeddings using a distance-based regression loss and bert," in *Proc. Int. Semantic Web Conf.*, 2023, pp. 561–578.
- [91] Y. He, Z. Yuan, J. Chen, and I. Horrocks, "Language models as hierarchy encoders," 2024, *arXiv:2401.11374*.
- [92] M. D. Ma, M. Chen, T.-L. Wu, and N. Peng, "HyperExpan: Taxonomy expansion with hyperbolic representation learning," in *Proc. Findings Conf. Empir. Methods Natural Lang. Process.*, 2021, pp. 4182–4194.
- [93] J. Hao et al., "MEDTO: Medical data to ontology matching using hybrid graph neural networks," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2021, pp. 2946–2954.
- [94] M. Nickel and D. Kiela, "Learning continuous hierarchies in the Lorentz model of hyperbolic geometry," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 3779–3788.
- [95] I. Chami, Z. Ying, C. Ré, and J. Leskovec, "Hyperbolic graph convolutional neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 4868–4879.
- [96] Z. Tang, T. Hinnerichs, X. Peng, X. Zhang, and R. Hoehndorf, "FALCON: Sound and complete neural semantic entailment over ALC ontologies," 2022. [Online]. Available: <https://arxiv.org/abs/2208.07628>
- [97] R. Abboud, I. Ceylan, T. Lukaszewicz, and T. Salvatori, "BoxE: A box embedding model for knowledge base completion," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 9649–9661.
- [98] Y. He, J. Chen, D. Antonyrajah, and I. Horrocks, "BERTMap: A BERT-based ontology alignment system," in *Proc. AAAI Conf. Artif. Intell.*, 2022, pp. 5684–5691.
- [99] W. Zhang et al., "Iteratively learning embeddings and rules for knowledge graph reasoning," in *Proc. World Wide Web Conf.*, 2019, pp. 2366–2377.
- [100] R. Xie et al., "Representation learning of knowledge graphs with hierarchical types," in *Proc. Int. Joint Conf. Artif. Intell.*, 2016, pp. 2965–2971.
- [101] L. Otero-Cerdeira, F. J. Rodríguez-Martínez, and A. Gómez-Rodríguez, "Ontology matching: A literature review," *Expert Syst. Appl.*, vol. 42, no. 2, pp. 949–971, 2015.
- [102] Y. He, J. Chen, H. Dong, E. Jiménez-Ruiz, A. Hadian, and I. Horrocks, "Machine learning-friendly biomedical datasets for equivalence and subsumption ontology matching," in *Proc. Int. Semantic Web Conf.*, 2022, pp. 575–591.
- [103] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian, and J. Lee, "Augmenting ontology alignment by semantic embedding and distant supervision," in *Proc. Int. Semantic Web Conf.*, 2021, pp. 392–408.
- [104] A. Ritchie, J. Chen, L. J. Castro, D. Rebbholz-Schuhmann, and E. Jiménez-Ruiz, "Ontology clustering with owl2vec," in *Proc. CEUR Workshop Proc.*, 2021, pp. 54–61.
- [105] J. Chen et al., "Zero-shot and few-shot learning with knowledge graphs: A comprehensive survey," *Proc. IEEE*, vol. 111, no. 6, pp. 653–685, Jun. 2023.
- [106] J. Chen, Y. Geng, Z. Chen, I. Horrocks, J. Z. Pan, and H. Chen, "Knowledge-aware zero-shot learning: Survey and perspective," in *Proc. Int. Joint Conf. Artif. Intell.*, 2021, pp. 4366–4373.
- [107] Y. Xian, B. Schiele, and Z. Akata, "Zero-shot learning—the good, the bad and the ugly," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4582–4591.
- [108] J. Chen, F. Lécué, Y. Geng, J. Z. Pan, and H. Chen, "Ontology-guided semantic composition for zero-shot learning," in *Proc. Int. Conf. Princ. Knowl. Representation Reasoning*, 2020, pp. 850–854.
- [109] P. N. Robinson, S. Köhler, S. Bauer, D. Seelow, D. Horn, and S. Mundlos, "The human phenotype ontology: A tool for annotating and analyzing human hereditary disease," *Amer. J. Hum. Genet.*, vol. 83, no. 5, pp. 610–615, Nov. 2008.
- [110] E. Choi, M. T. Bahadori, L. Song, W. F. Stewart, and J. Sun, "Gram: Graph-based attention model for healthcare representation learning," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 787–795.
- [111] X. Peng, G. Long, T. Shen, S. Wang, and J. Jiang, "Sequential diagnosis prediction with transformer and ontological representation," in *Proc. IEEE Int. Conf. Data Mining*, 2021, pp. 489–498.
- [112] K. Niu, Y. Lu, X. Peng, and J. Zeng, "Fusion of sequential visits and medical ontology for mortality prediction," *J. Biomed. Inform.*, vol. 127, 2022, Art. no. 104012.
- [113] C. W. Cheong, K. Yin, W. K. Cheung, B. C. Fung, and J. Poon, "Adaptive integration of categorical and multi-relational ontologies with EHR data for medical concept embedding," *ACM Trans. Intell. Syst. Technol.*, vol. 14, no. 6, pp. 1–20, 2023.
- [114] M. Zhang, C. R. King, M. Avidan, and Y. Chen, "Hierarchical attention propagation for healthcare representation learning," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 249–256.
- [115] L. Song, C. W. Cheong, K. Yin, W. K. Cheung, B. C. Fung, and J. Poon, "Medical concept embedding with multiple ontological representations," in *Proc. Int. Joint Conf. Artif. Intell.*, 2019, pp. 4613–4619.
- [116] F. Ma, Q. You, H. Xiao, R. Chitta, J. Zhou, and J. Gao, "KAME: Knowledge-based attention model for diagnosis prediction in healthcare," in *Proc. ACM Int. Conf. Inf. Knowl. Manage.*, 2018, pp. 743–752.
- [117] Z. Yao, B. Liu, F. Wang, D. Sow, and Y. Li, "Ontology-aware prescription recommendation in treatment pathways using multi-evidence healthcare data," *ACM Trans. Inf. Syst.*, vol. 41, no. 4, pp. 1–29, 2023.
- [118] F. Shen et al., "HPO2Vec: Leveraging heterogeneous knowledge resources to enrich node embeddings for the human phenotype ontology," *J. Biomed. Inform.*, vol. 96, 2019, Art. no. 103246.
- [119] S. Nunes, R. T. Sousa, and C. Pesquita, "Multi-domain knowledge graph embeddings for gene-disease association prediction," *J. Biomed. Semantics*, vol. 14, no. 1, Aug. 2023, Art. no. 11.
- [120] A. Althagafi, F. Zhapa-Camacho, and R. Hoehndorf, "Prioritizing genomic variants through neuro-symbolic, knowledge-enhanced learning," *Bioinformatics*, vol. 40, 2024, Art. no. btac 301.
- [121] K. Agarwal, T. Eftimov, R. Addanki, S. Choudhury, S. Tamang, and R. Rallo, "Snomed2vec: Random walk and poincaré embeddings of a clinical knowledge base for healthcare analytics," 2019, *arXiv: 1907.08650*.
- [122] J. Vilela et al., "Biomedical knowledge graph embeddings for personalized medicine: Predicting disease-gene associations," *Expert Syst.*, vol. 40, no. 5, Nov. 2022, Art. no. e13181.
- [123] Y. Wang, P. Wegner, D. Domingo-Fernández, and A. Tom Kodamullil, "Multi-ontology embeddings approach on human-aligned multi-ontologies representation for gene-disease associations prediction," *Heliyon*, vol. 9, no. 11, Nov. 2023, Art. no. e21502.
- [124] C. Zhao, T. Liu, and Z. Wang, "PANDA2: Protein function prediction using graph neural networks," *NAR Genomic. Bioinf.*, vol. 4, no. 1, 2022, Art. no. lqac004.
- [125] S. Qiu, G. Yu, X. Lu, C. Domeniconi, and M. Guo, "Isoform function prediction by gene ontology embedding," *Bioinformatics*, vol. 38, no. 19, pp. 4581–4588, 2022.
- [126] M. Kulmanov, F. J. Guzmán-Vega, P. Duek Roggli, L. Lane, S. T. Arold, and R. Hoehndorf, "Protein function prediction as approximate semantic entailment," *Nature Mach. Intell.*, vol. 6, no. 2, pp. 220–228, Feb. 2024.
- [127] W. Li et al., "Partial order relation-based gene ontology embedding improves protein function prediction," *Brief. Bioinf.*, vol. 25, no. 2, 2024, Art. no. bbae077.
- [128] Z. Lin et al., "Evolutionary-scale prediction of atomic-level protein structure with a language model," *Science*, vol. 379, no. 6637, pp. 1123–1130, 2023.

- [129] A. Elnaggar et al., "ProtTrans: Toward understanding the language of life through self-supervised learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 7112–7127, Oct. 2022.
- [130] A. A. Salatino, F. Osborne, T. Thanapalasingam, and E. Motta, "The cso classifier: Ontology-driven detection of research topics in scholarly articles," in *Proc. 23rd Int. Conf. Theory Pract. Digit. Libraries Digit. Libraries Open Knowl.*, Springer, 2019, pp. 296–311.
- [131] F. Ali, S. El-Sappagh, and D. Kwak, "Fuzzy ontology and LSTM-based text mining: A transportation network monitoring system for assisting travel," *Sensors*, vol. 19, no. 2, 2019, Art. no. 234.
- [132] A. H. Sweidan, N. El-Bendary, and H. Al-Feel, "Sentence-level aspect-based sentiment analysis for classifying adverse drug reactions (ADRs) using hybrid ontology-XLNet transfer learning," *IEEE Access*, vol. 9, pp. 90828–90846, 2021.
- [133] S. Deng et al., "OntoED: Low-resource event detection with ontology embedding," 2021, *arXiv:2105.10922*.
- [134] C. Erten and D. Kazakov, "Ontology graph embeddings and ILP for financial forecasting," in *Proc. 30th Int. Conf. Inductive Log. Program.*, Springer, 2021, pp. 111–124.
- [135] F. Dassereto, L. Di Rocco, G. Guerrini, and M. Bertolotto, "Evaluating the effectiveness of embeddings in representing the structure of geospatial ontologies," in *Proc. AGILE Conf. Geographic Inf. Sci.*, Springer, 2020, pp. 41–57.
- [136] M. Wick, B. Vatan, and B. Christophe, "Geonames ontology," 2015. [Online]. Available: <http://www.geonames.org/ontology>
- [137] M. Horridge and S. Bechhofer, "The OWL API: A java API for OWL ontologies," *Semant. Web*, vol. 2, no. 1, pp. 11–21, Jan. 2011.
- [138] M. van Bekkum, M. de Boer, F. van Harmelen, A. Meyer-Vitali, and A. T. Teije, "Modular design patterns for hybrid learning and reasoning systems: A taxonomy, patterns and use cases," *Appl. Intell.*, vol. 51, no. 9, pp. 6528–6546, Jun. 2021.
- [139] Y. Chang et al., "A survey on evaluation of large language models," *ACM Trans. Intell. Syst. Technol.*, vol. 15, no. 3, pp. 1–45, 2024.
- [140] H. Touvron et al., "LLaMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.
- [141] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, "Unifying large language models and knowledge graphs: A roadmap," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 7, pp. 3580–3599, Jul. 2024.
- [142] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 9459–9474.
- [143] P. Hitzler, F. Bianchi, M. Ebrahimi, and M. K. Sarker, "Neural-symbolic integration and the Semantic Web," *Semantic Web*, vol. 11, no. 1, pp. 3–11, 2020.
- [144] Y. Bian and X.-Q. Xie, "Generative chemistry: Drug discovery with deep learning generative models," *J. Mol. Model.*, vol. 27, pp. 1–18, 2021.
- [145] M. Jullien, M. Valentino, H. Frost, P. O'Regan, D. Landers, and A. Freitas, "Semeval-2023 task 7: Multi-evidence natural language inference for clinical trial data," 2023, *arXiv:2305.02993*.



Jiaoyan Chen is a lecturer with the Department of Computer Science, University of Manchester and a part-time senior researcher with the Department of Computer Science, University of Oxford. His main interests include KG, ontology and machine learning.



Olga Mashkova is currently working toward the PhD degree in computer science with the King Abdullah University of Science and Technology. Her main interests include deep learning, bioinformatics and knowledge representation and reasoning.



Fernando Zhapa-Camacho is currently working toward the PhD degree in computer science with the King Abdullah University of Science and Technology. His main interests include knowledge representation and reasoning, machine learning and bioinformatics.



Robert Hoehndorf is an associate professor in computer science with the King Abdullah University of Science and Technology. His main interests include artificial intelligence, knowledge representation, biomedical informatics and ontology.



Yuan He is a researcher with the Department of Computer Science, University of Oxford. His main interests include large language models, knowledge engineering and neural-symbolic integration.



Ian Horrocks is a professor in computer science with the University of Oxford. His main interests include knowledge representation, ontologies and ontology languages, description logics, and the Semantic Web.