

Metapath-Based Data-Augmentation for Knowledge Graphs

Saurav Manchanda
manch043@umn.edu

ABSTRACT

Knowledge graph (KG) embedding techniques use structured relationships between entities to learn low-dimensional representations of entities and relations. The traditional KG embedding techniques (such as TransE and DistMult) estimate these embeddings via simple models developed over observed KG triplets. These approaches differ in their triplet scoring loss functions. As these models only use the observed triplets to estimate the embeddings, they are prone to suffer through data sparsity that usually occurs in the real-world knowledge graphs, i.e., the lack of enough triplets per entity. In this paper, we propose an efficient method to augment the number of triplets to address the problem of data sparsity. We use random walks to create additional triplets, such that the relations carried by these introduced triplets correspond to the metapath (composition of the sequence of underlying relations) induced by the random walks. We also provide approaches to accurately and efficiently choose the informative metapaths from the possible set of metapaths, generated by the random walks. The proposed approaches are model-agnostic, and the augmented training dataset can be used with any KG embedding approach out of the box. Experimental results obtained on the benchmark datasets show the advantages of the proposed approach.

KEYWORDS

knowledge graph embedding, knowledge graph augmentation, random walk, leaning on graphs

1 INTRODUCTION

Knowledge graphs (KGs) are data structures that store information about different entities (nodes) and their relations (edges). They are used to organize information in many domains such as music [8], movies [26], (e-)commerce [38], and sciences [15]. A common approach to using KGs in various information retrieval and machine learning tasks is to compute knowledge graph embeddings (KGE) [11, 31]. These approaches embed a KG’s entities and relation types associated with an edge into a d -dimensional space such that the embedding vectors associated with the entities and the relation types satisfy a pre-determined mathematical model. Numerous models for computing knowledge graph embeddings have been developed, such as TransE [5], TransR [22], and DistMult [37].

There are many possible facts and relations in the real world, but collecting them can be challenging; consequently, existing knowledge graphs lack completeness. Learning embeddings for entities and relations in knowledge graphs can be considered knowledge induction, and the induced embeddings can be used to infer new triplets in the knowledge graph. The embedding quality is heavily reliant on the quantity and quality of the underlying triplets. This makes obtaining meaningful embeddings from sparse KG difficult, because of the lack of enough number of triplets per entity or relation. In addition, KG embedding approaches that use negative sampling encounter an increased false negative rate in the chosen

negative samples when dealing with sparse KGs. Negative sampling is performed by creating synthetic observations (relating two entities using a relation) such that the created observation is not observed in the training data. The chances of selecting an actual positive observation as a negative sample are higher with a sparse KG.

One solution to address the sparsity of triplets is to perform data augmentation. Various forms of data augmentation are possible. For example, prior approaches have explored augmenting the KG triplets with text corpus [14, 33], using schema and semantic constraints [18, 35], using inference rules on the KG [13, 14, 39, 40], etc. The approaches proposed in this paper resemble the ones that leverage the KG inference rules (multi-hop reasoning, in particular) to augment the KG triplets. Specifically, multi-hop reasoning is the task of learning explicit inference patterns in a KG. For example, if the KG includes the facts such as b is a ’s mother, c is b ’s husband, we want to learn the following pattern: LHS: $\text{isMother}(a, b) \ \& \ \text{isHusband}(b, c) \implies$ RHS: $\text{isFather}(a, c)$. Various approaches leverage such reasoning to expand the set of KG triplets, and hence, learn higher-quality embeddings. The question we answer is: *what if the isFather relation (RHS in the above pattern) is not present in the KG? Could we still leverage the composite relation isMother(a, b) & isHusband(b, c) (LHS in the above pattern) to augment the triplets?* This paper proposes a solution in this direction. Our proposed approach performs data augmentation predicated on composite relations, even if the composite relations could not be mapped to existing relations.

Specifically, we provide approaches that use random walks to efficiently and accurately estimate *informative* (multi-hop) composite relations (also called metapaths) in the KG. Each metapath can be effectively considered a new relation in the KG, thus augmenting the number of triplets in the KG. However, the difficulty in considering each metapath as a separate relation lies in the exponential size of the search space: the number of possible metapaths increases as we increase the length of the metapath. To this end, we also provide various approaches that allow parameter sharing among the relations, thus alleviating the curse of dimensionality incurred because of the exponential search space. In summary, the main contributions are as follows:

- We propose a quantitative way to measure the *metapath information*, i.e., the information encoded by a metapath.
- We introduce computationally efficient and accurate approaches to use random walks to estimate *informative* metapaths in the KG and use them to augment the KG triplets.
- We provide various parameter-sharing approaches to tackle the curse of dimensionality induced by the exponential search space of the possible metapaths.
- We experimentally evaluate the performance of the proposed framework on different knowledge graph benchmarks. The proposed approach outperforms the baselines.

2 RELATED WORK

The research areas relevant to the work present in this paper belong to *Knowledge Graph Embeddings*, *Inference Rule Mining*, and *Data Augmentation for KGE* approaches. We discuss these areas below:

2.1 Knowledge Graph Embeddings

Knowledge graph embedding (KGE) is a machine learning task of learning a low-dimensional representation of a knowledge graph’s entities and relations while preserving their semantic meaning. Leveraging their embedded representation, knowledge graphs (KGs) can be used for various applications such as link prediction, triple classification, entity recognition, clustering, and relation extraction [1, 16]. KGE differs from ordinary relation inference as the information in a knowledge graph is multi-relational and more complex to model and computationally expensive. Many models have been devoted to estimating KGE. Some translation-based embedding models, such as TransE [5], TransH [34], TransR [22] and TransD [16] apply simple linear or bilinear operations to model the embeddings of entities and relations. DistMult [37] and ComplEx [30] design similarity scoring functions to learn semantic information. ConvE [7] and ConvKB [25] apply a convolutional neural network to learn non-linear features.

In addition, there are approaches that tackle representation learning on homogeneous graphs, i.e., graphs with just one type of relation. Examples include DeepWalk [27] and node2vec [12]. These approaches are based on random walks for embedding generation. For each node in the graph, these models generate a random path of connected nodes. The random path gives a sequence of nodes, and we can train models like Word2Vec (skip-gram) [23] to obtain the node embeddings. The approaches presented in this paper are motivated by these random-walk-based embedding generation approaches.

2.2 Rule mining

Rule mining in KGs finds rules such as $\text{isMarried}(a,b) \implies \text{staysTogether}(a,b)$. Approaches like Inductive Logic Programming (ILP) and AMIE [10] are popular choices for learning such rules. AMIE is an improved version of a fuzzy Horn rule-mining algorithm based on inductive logic programming (ILP). It overcomes the difficulty of traditional rule-mining algorithms in generating negative evidence on the basis of the open-world assumption. However, the difficulty in finding such rules lies in the exponential size of the search space: every relation can potentially be combined with every other relation in a rule. To tackle this, several extensions of the AMIE have been proposed [9, 19].

One particular type of rule mining finds its use in multi-hop reasoning. Multi-hop reasoning is the task of learning explicit inference patterns in a KG. For example, if the KG includes the facts such as b is a ’s mother, c is b ’s husband, we want to learn the following rule (pattern): $\text{isMother}(a, b) \ \& \ \text{isHusband}(b, c) \implies \text{isFather}(a, c)$. Path-Ranking Algorithm (PRA) [20] is a popular algorithm for estimating such patterns in KGs. Lao et al. [20] uses a soft inference procedure based on a combination of constrained, weighted, random walks through the knowledge base graph to infer rules for the KG. Their solution can learn to infer different target relations by tuning the weights associated with random walks that follow

different paths through the graph, using a version of the PRA. Wang and Li [32] proposes a rule learning approach RDF2Rules for RDF KGs. Rules are learned by finding frequent predicate cycles in RDF graphs. A new confidence measure is also proposed for evaluating the reliability of the mined rules. Barati et al. [3] intro an approach called SWARM (Semantic Web Association Rule Mining) to mine semantic association rules from RDF KG. SWARM leverages the schema information to enrich the semantics of the rules. [36] proposes a reinforcement learning framework to find reasoning paths in the KG. Unlike random walk-based models, the RL model allows controlling the properties of the found paths. These effective paths can also be used as an alternative to PRA in many path-based reasoning methods.

2.3 Data augmentation for KGE

Various approaches have been developed to augment the KG data. Wang et al. [33] augments the KG data with a text corpus. Specifically, they introduce additional text-based information and performed joint model training by combining triplets from a knowledge graph with entity descriptions in text. Guo et al. [13] proposed a method that embeds rules and triplets jointly. This method learns rules based on the TransE model and retrains the embeddings combined with rules. Zhang et al. [39] proposed IterE, to alleviate the sparsity of entities in KG. In this framework, embeddings and rules are learned iteratively; the embeddings are learned from existing triplets, and new triplets are inferred from the rules that have been learned. Zhao et al. [40] iteratively trains embeddings and discovers patterns between relations to construct new triplets. For example, the model can discover the pattern that *located_in* is the inverse relation of *has_city*. If the KG contains the triplet (*seattle, located_in, usa*), then according to the discovered rules, the model can infer another triplet, (*usa, has_city, seattle*). Ho et al. [14] iteratively extends rules induced from a KG by relying on feedback from a precomputed embedding model over the KG and external information sources including text corpora. Xie et al. [35] adds entity-level information into the representation learning model by using an encoder. This approach can be used to create fused semantic vectors for entities at different levels to improve the representation effect. Krompaß et al. [18] adds prior knowledge about the semantics of relation-types, extracted from the schema of the knowledge graph (type-constraints) or approximated through a local closed-world assumption, to the representation learning model. Only triplets that satisfy the type constraints are considered for modeling; thus, irrelevant or erroneous samples can be effectively filtered out to a certain extent.

3 DEFINITIONS AND NOTATIONS

A graph is composed of nodes and edges $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of edges. A knowledge graph (KG) is a special type of graph whose nodes and edges have types. A node in a knowledge graph represents an entity and an edge represents a relation between two entities. The edges are in the form of triplets (h, r, t) , which indicates that a pair of entities h (head) and t (tail) are connected to each other via a relation (edge-type) r . Each node $v \in \mathcal{V}$ and each edge $e \in \mathcal{E}$ are associated with their type mapping functions $\tau(v) : \mathcal{V} \rightarrow \mathcal{A}$ and $\phi(e) : \mathcal{E} \rightarrow \mathcal{R}$, respectively. A

Table 1: Notation used throughout the paper.

Symbol	Description
\mathcal{G}	A (knowledge) graph.
\mathcal{V}	Set of nodes or entities in the graph \mathcal{G} .
\mathcal{E}	Set of edges in the graph \mathcal{G} .
\mathcal{A}	Set of node types in the graph \mathcal{G} .
\mathcal{R}	Set of edge types in the graph \mathcal{G} .
\mathfrak{m}	A metapath in the graph \mathcal{G} , having the edge-type sequence $[\mathfrak{m}_1, \mathfrak{m}_2, \dots, \mathfrak{m}_{ \mathfrak{m} }]$, where $ \mathfrak{m} $ is the length of the metapath.
$\tau(v)$	Type mapping function $\tau(v) : \mathcal{V} \rightarrow \mathcal{A}$, that maps a node to its corresponding type in the graph \mathcal{G} .
$\phi(e)$	Type mapping function $\tau(e) : \mathcal{E} \rightarrow \mathcal{R}$, that maps an edge to its corresponding type in the graph \mathcal{G} .
$n_e^{\mathfrak{m}}(\mathcal{G})$	Number of distinct metapaths of type \mathfrak{m} that the edge e is a part of in the graph \mathcal{G} .
$assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}}$	Association measure, defined as the fraction of edges of the type \mathfrak{m}_i that are a part of the metapath instance for the metapath \mathfrak{m} .
$z_{\mathfrak{m}}$	Metapath information for the metapath \mathfrak{m} .
$conf_{\mathfrak{m} \rightarrow q}$	Rule-confidence for the rule $\mathfrak{m} \rightarrow q$, defined as the fraction of node pairs (h, t) such that the triplet (h, q, t) exists given that the path (h, \mathfrak{m}, t) exists.
r_v	Representation (embedding) of v where v is an entity, relation or basis vector.

metapath \mathfrak{m} on the knowledge graph \mathcal{G} is defined as a sequence of relations (edge-types) $\mathfrak{m} = [\mathfrak{m}_1, \mathfrak{m}_2, \dots, \mathfrak{m}_{|\mathfrak{m}|}]$, which describes a composite relation $\mathfrak{m}_1 \circ \mathfrak{m}_2 \circ \dots \circ \mathfrak{m}_{|\mathfrak{m}|}$. Given a metapath \mathfrak{m} , a metapath instance p of \mathfrak{m} is defined as a node and edge sequence in the graph following the sequence of relations defined by \mathfrak{m} . The notation used in this paper is summarized in Table 1.

4 PROPOSED APPROACH: WALK-AND-RELATE (WAR)

Walk-and-Relate (WAR) uses short random walks to construct sequences of connected nodes in the KG. Similar to Deepwalk [27], the random walk samples uniformly from the neighbors of the last vertex visited until the maximum length (l_{max}) is reached¹. For each node in the graph, WAR performs a random walk, thus sampling a random path of nodes connected to it. For each pair of nodes in this random walk, we can effectively consider them *neighbors* in the graph, associated by a metapath (composite relation). Thus, we use these pairs of nodes, along with the associated metapath to construct additional triplets and augment the training dataset for KG embeddings. While the nodes in the augmented triplets can be directly used by the triplet scoring function of the desired KG embedding method, there are some unknowns regarding how to use the metapaths. Particularly:

- *Does the composite relation described by the metapath provide any additional information?* Augmenting with *non-informative* metapaths could bring undesired noise to the training triplets. This is expected to be the case as the number of possible metapaths grows exponentially with the length of the random walk. To address this, we take motivation from the vast literature on

¹In principle, there is no requirement for the random walks to be of the same length and we can have restarts in the random walks

association mining and provide approaches to efficiently mine *informative* metapaths and quantify the information provided by them (details in Section 4.1).

- *How do we represent the metapath as a single relation which is a requirement for KG embedding models based on triplet scoring?* For each selected informative metapath (\mathfrak{m}), we check if the metapath can be confidently mapped to one of the existing relations in the KG (rule-mining: $\mathfrak{m} \implies r$). If yes, we use the corresponding rule to generate new triplets for the relation r (Note that this step is similar to earlier approaches). Details on rule-mining are covered in Section 4.2. If the metapath cannot be confidently mapped to one of the existing relations, each metapath is effectively considered a new relation in the KG. However, the difficulty in considering each metapath as a separate relation lies in the exponential size of the search space. To address this curse of dimensionality, we provide multiple solutions based on generalization via parameter sharing (details in Section 4.3).

The next sections cover the WAR framework in detail.

4.1 How informative is a metapath?

We borrow ideas from the association rule mining literature [2] to measure the information conveyed through a metapath, which we call metapath-information. First, we define an association measure, denoted by $assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}}$, that measures how often an edge with type \mathfrak{m}_i is part of a metapath instance of type \mathfrak{m} . Similar to the confidence measure used in the association rule mining, we define the $assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}}$ as the fraction of edges of the type \mathfrak{m}_i that are a part of the metapath instance for the metapath \mathfrak{m} . Formally, the definition is given as

$$assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}} = \frac{\sum_{e \in \mathcal{G}: \phi(e) = \mathfrak{m}_i} \mathbb{1}(n_e^{\mathfrak{m}}(\mathcal{G}) > 0)}{\sum_{e \in \mathcal{G}: \phi(e) = \mathfrak{m}_i} \mathbb{1}}, \quad (1)$$

where $n_e^{\mathfrak{m}}(\mathcal{G})$ denotes the distinct metapath instances of type \mathfrak{m} that the edge e is a part of in the graph \mathcal{G} . Next, we define the metapath-information of the metapath \mathfrak{m} as the product of the association measures $assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}}$ of each of its constituent relations $\mathfrak{m}_i : 1 \leq i \leq |\mathfrak{m}|$. Formally, we denote the metapath-information of \mathfrak{m} as $z_{\mathfrak{m}}$ and calculate it as

$$z_{\mathfrak{m}} = \prod_{i=1}^{|\mathfrak{m}|} assoc_{\mathfrak{m}_i \rightarrow \mathfrak{m}}.$$

Estimation of $z_{\mathfrak{m}}$ requires enumeration of all the metapath instances of \mathfrak{m} . Trivial enumerating these metapath instances would lead to a time and space complexity of the order of $O(|\mathcal{E}|^{|\mathfrak{m}|})$. The exponential complexity raises scalability concerns and prohibits estimating the metapath importance for graphs with a large number of edges.

In this direction, we provide two heuristics to address the scalability problems. First, we provide an approach to prune the exploration space of the candidate metapaths, leveraging the downward-closure property of the metapath-information. Second, we provide an approach to estimate the metapath-information on a sampled graph.

4.1.1 Pruning the exploration space. Our goal is to select the informative metapaths, and not to calculate the metapath-information $z_{\mathfrak{m}}$ of all possible metapaths. Thus, we employ a pruning strategy, i.e., do not compute $z_{\mathfrak{m}}$ for a metapath \mathfrak{m} if we have *evidence* that

z_m is less than a threshold. We leverage the downward-closure property of the metapath-information to compute this *evidence*.

LEMMA 4.1. *Metapath-information is downward-closed (anti monotonic), i.e., if $m = [m_1, m_2, \dots, m_{|m|}]$ and $\hat{m} = [m_1, m_2, \dots, m_{|m|-1}]$, then $z_m \leq z_{\hat{m}}$.*

Proof: Each metapath instance of type m can be mapped to a metapath instance of type \hat{m} , by removing the edge of relation $m_{|m|}$. Thus,

$$\begin{aligned} n_e^{m_i} &\leq n_e^{\hat{m}_i}, \forall 1 \leq i \leq |\hat{m}| (= |m| - 1) \\ \implies \text{assoc}_{m_i \rightarrow m} &\leq \text{assoc}_{\hat{m}_i \rightarrow \hat{m}}, \forall 1 \leq i \leq |\hat{m}| \end{aligned}$$

Next, the metapath-information z_m is expressed as

$$\begin{aligned} z_m &= \prod_{i=1}^{|m|} \text{assoc}_{m_i \rightarrow m} \leq \prod_{i=1}^{|\hat{m}|} \text{assoc}_{\hat{m}_i \rightarrow \hat{m}} \times \text{assoc}_{m_{|m|} \rightarrow m} \\ \implies z_m &\leq z_{\hat{m}} \times \text{assoc}_{m_{|m|} \rightarrow m} \end{aligned}$$

As a fraction, the association $\text{assoc}_{m_i \rightarrow m} \leq 1$. Thus, $z_m \leq z_{\hat{m}}$.

Following the above downward-closure property, we only need to expand the exploration space to consider the metapath $m = [m_1, m_2, \dots, m_{|m|}]$ only if the metapath $\hat{m} = [m_1, m_2, \dots, m_{|m|-1}]$ is informative ($z_{\hat{m}}$ is above a threshold).

4.1.2 *Estimation on a sampled graph.* Instead of estimation on the full graph, we sample the edges from the graph and estimate the metapath-information from there. However, metapath-information estimated from the sampled graph (with randomly sampled edges) is lower (thus, not representative) than the metapath-information on the full graph². This is because the numerator in Equation 1 decays faster with sampling as compared to the denominator. To address this difference, we introduce a correction factor to estimate the metapath-information on the full KG, using the sampled KG.

Given an edge e of type m_i , it does not belong to any metapath instance of type m (i.e., $n_e^m(\mathcal{G}_S) = 0$) if one of the following conditions hold:

- e does not belong to any metapath instance of type m in the original unsampled graph, i.e., $n_e^m(\mathcal{G}) = 0$. The probability of this condition holding is

$$p \times \mathbb{1}(n_e^m(\mathcal{G}) = 0),$$

where p denotes the sampling probability. i.e., $p = P(e \in \mathcal{G}_S)$

- e belongs to one or more metapath instances of type m in the original unsampled (full) graph, i.e., $n_e^m(\mathcal{G}) > 0$, but each of those metapath instances does not appear in the sampled graph. This is possible if at least one of the constituent edges of those metapath instances (apart from e) is removed because of sampling. The probability of this condition holding is

$$\underbrace{p}_{e \text{ is retained}} \times \underbrace{n_e^m(\mathcal{G})^{(1-p)^{|m|-1}}}_{\text{at least one edge is removed}} \times \mathbb{1}(n_e^m(\mathcal{G}) > 0)$$

Combining the above two conditions, and aggregating over the observed number of edges of type m_i in the sampled graph (\mathcal{G}_S) that do not belong to any metapath instance of type m , we get:

²In principle, the notion of missing edges in a KG means that KG itself is sampled from the *complete* KG, which we don't have access to. Thus, the metapath-information estimated from a given KG is a lower bound on the "ideal" metapath-information.

$$\begin{aligned} \sum_{e \in \mathcal{G}_S: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}_S) = 0) &= p \sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}) = 0) \\ &+ p \sum_{e \in \mathcal{G}: \phi(e)=m_i} (1-p)^{|m|-1} n_e^m(\mathcal{G}) \quad (2) \end{aligned}$$

To simplify the above formulation, we make the following observations:

- Assuming $n_e^m(\mathcal{G})$ is uniformly distributed³ across the edges, we get

$$n_e^m(\mathcal{G}) \approx \frac{n^m(\mathcal{G})}{\sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}) > 0)},$$

where $n^m(\mathcal{G})$ is the number of metapaths instances of type m in the graph \mathcal{G} . Further, for a metapath instance present in the unsampled graph, the probability of it being present in the sampled graph is $p^{|m|}$. Thus, $p^{|m|} \times n^m(\mathcal{G}) \approx n^m(\mathcal{G}_S)$.

- For any metapath instance present in the unsampled graph, the probability of it being present in the sampled graph is $p^{|m|}$. Thus, $p^{|m|} \times n^m(\mathcal{G}) \approx n^m(\mathcal{G}_S)$.
- Being a part of a metapath instance or not is both mutually exclusive and exhaustive. Thus,

$$\begin{aligned} \sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}) = 0) + \sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}) > 0) \\ = \sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1} \end{aligned}$$

Substituting the above observations in Equation 2 gives us

$$\begin{aligned} p \left[\sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1} - x + x(1-p)^{|m|-1} \frac{n^m(\mathcal{G}_S)}{p^{|m|x}} \right] \\ - \sum_{e \in \mathcal{G}_S: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}_S) = 0) = 0, \quad (3) \end{aligned}$$

where $x = \sum_{e \in \mathcal{G}: \phi(e)=m_i} \mathbb{1}(n_e^m(\mathcal{G}) > 0)$. Except x , all other variables are observed or hyperparameters (p) in Equation 3. Thus, various root-finding algorithms could be used to find the optimal value of x that satisfy Equation 3, such as Brent's [6] and Golden-section search [17]. The solution x is then substituted for the numerator in the Equation 1 to get the corrected association-measure $\text{assoc}_{m_i \rightarrow m}$, and thus accurately estimate the metapath-information z_m .

Algorithm 1 provides the pseudo-code for estimating the informative metapaths.

4.2 How to map the informative metapath to relations in the graph?

For the augmentation triplets (h, m, t) to be used by the desired KG embedding approach that uses triplet scoring functions, we represent the metapath m as a single relation in the KG. To address this, we mine rules of the form $m \rightarrow q$, using the standard notion of confidence used in the association rule mining. Specifically, we define the rule confidence (denoted by $\text{conf}_{m \rightarrow q}$) for the rule $m \rightarrow q$

³The assumption gives a lower bound (following the AM-GM inequality property)

Algorithm 1 Estimating metapath-information

Input: Graph \mathcal{G} , maximum length of metapaths to discover l_{max} , metapath-information threshold t , sampling probability p .
Output: Metapaths with associated metapath-information: result_set

- 1: Sample p edges randomly from \mathcal{G} as \mathcal{G}_S
- 2: Represent \mathcal{G}_S in relational table format (T), with three columns: source (src), relation (rel), and destination (dst).
- 3: result_set = {}, $\hat{T} = T$
- 4: **for** $l \in [2, l_{max}]$ **do**
- 5: $\hat{T} = \hat{T}$ inner join T on $\hat{T}.dst = T.src$
- 6: **for** each metapath \mathbf{m} , group g in \hat{T} group by rel columns **do**
- 7: Initialize $z_{\mathbf{m}} = 1$
- 8: **for** each relation $m_i \in \mathbf{m}$ **do**
- 9: Calculate the association $assoc_{m_i \rightarrow \mathbf{m}}$ with correction factor using Equations 1 and 3.
- 10: Update $z_{\mathbf{m}} = z_{\mathbf{m}} * assoc_{m_i \rightarrow \mathbf{m}}$.
- 11: **if** $z_{\mathbf{m}} < t$ **then**
- 12: Remove the instances corresponding to \mathbf{m} from \hat{T} , i.e., $\hat{T} = \hat{T} - g$
- 13: **break**
- 14: **end if**
- 15: **end for**
- 16: **if** $z_{\mathbf{m}} \geq t$ **then**
- 17: Add $\{\mathbf{m} : z_{\mathbf{m}}\}$ to the result_set.
- 18: **end if**
- 19: **end for**
- 20: **end for**
- 21: **return** result_set

as the fraction of node pairs (h, t) such that the triplet (h, q, t) exists given that the path (h, \mathbf{m}, t) exists. Formally,

$$conf_{\mathbf{m} \rightarrow q} = \frac{\sum_{h,t \in \mathcal{V}} \mathbb{1}((h, q, t) \in \mathcal{G} \wedge (h, \mathbf{m}, t) \in \mathcal{G})}{\sum_{h,t \in \mathcal{V}} \mathbb{1}((h, \mathbf{m}, t) \in \mathcal{G})}.$$

For a metapath \mathbf{m} , we denote $rulemap(\mathbf{m})$ as the mapping that holds the key-value pairs $[q, conf_{\mathbf{m} \rightarrow q}]$. Since we are only interested in high-confidence rules, we have a minimum threshold on the $conf_{\mathbf{m} \rightarrow q}$ for $[q, conf_{\mathbf{m} \rightarrow q}]$ to be added to $rulemap(\mathbf{m})$ (say, 0.5).

Next, for each informative-metapath \mathbf{m} , we map it to an individual relation with the following setup:

- (i) If the $rulemap(\mathbf{m})$ is non-empty, we sample a relation q out of $rulemap(\mathbf{m})$, with the sampling probability $conf_{\mathbf{m} \rightarrow q}$. The sampling is done independently for each augmentation triplet (h, \mathbf{m}, t) on the course of the training.
- (ii) If the $rulemap(\mathbf{m})$ is empty, we simply introduce a new relation for the metapath \mathbf{m} in the KG.

4.3 Generalization with parameter sharing

If a metapath cannot be mapped to existing relations in the KG, we introduce a new relation describing that metapath in the KG. The increase in the number of unique relations would lead to rapid growth in the number of parameters (relation embeddings). The number of parameters rises exponentially (in terms of the number of original relations in the KG) with the lengths of the considered

metapaths. This can easily lead to over-fitting on rare metapaths (which counters our primary goal of addressing data sparsity). To address this problem, we propose various approaches to allow parameter sharing among the embeddings for the metapaths, thus limiting the number of unique parameters and allowing for better generalization.

4.3.1 Using composition properties of the underlying model. Models such as TransE [5] and RotatE [29] have triplet scoring functions that enable them to infer the composition patterns. For example, given the metapath $\mathbf{m} = [m_1, m_2, \dots, m_{|\mathbf{m}|}]$, its representation $r_{\mathbf{m}}$ is computed as $r_{\mathbf{m}} = r_{m_1} + r_{m_2} + \dots + r_{m_{|\mathbf{m}|}}$ under the assumptions of the TransE model. Similarly, the representation $r_{\mathbf{m}}$ under the assumptions of the RotatE model is given as $r_{\mathbf{m}} = r_{m_1} \circ r_{m_2} \circ \dots \circ r_{m_{|\mathbf{m}|}}$, where \circ is the Hadamard (or element-wise product). Thus, we simply use the representations of the constituent relations to infer the representation of the new relation describing the metapath, without introducing any additional parameters. Note that this approach can only be applied to models that allow composition pattern inference.

4.3.2 Using parametric sequence models to represent the metapath. Parametric sequence models such as Recurrent Neural Networks (RNN) are employed to generate the representation of the metapath by using the sequence of representations of the constituent relations. The only additional parameters that are introduced by this approach correspond to the sequence model (RNN), which is a constant overhead. The RNN parameters are estimated using back-propagation from the loss function of the KGE model.

4.3.3 Using basis decomposition. Similar to Relational Graph Convolutional Network (R-GCN) [28], we use basis-decomposition as a way to share parameters among the representations for relations (and metapaths). With basis decomposition, the representation $r_{\mathbf{m}}$ is defined as follows:

$$r_{\mathbf{m}} = \sum_{b=1}^B a_{\mathbf{m},b} r_b,$$

i.e. we calculate the representation $r_{\mathbf{m}}$ as a linear combination of basis transformations $r_b : 1 \leq b \leq B$ with coefficients $a_{\mathbf{m},b} : 1 \leq b \leq B$ such that only the coefficients depend on \mathbf{m} . This approach does introduce additional parameters that grow exponentially with the lengths of the considered metapaths, but their number is much smaller as compared to the representation size.

4.4 Putting it all together

Given the set of informative metapaths and associated rulemaps, we use random walks to generate instances of those metapaths on the graph. These instances are used to create augmentation triplets in the KG and are directly taken as input by any downstream KG-embedding model based on triplet scoring loss functions. The KG-embedding model, thus, has access to a larger pool of training data, effectively leading to higher-quality KG embeddings. For the KGE model to be aware of the metapath-information and rule confidence, we use them as the edge weight for the KGE model (edge weights are used to scale the loss function for the corresponding triplets). Specifically, if we introduce a new relation for the metapath \mathbf{m} , we use $z_{\mathbf{m}}$ as the edge weight for the augmented triplets added as a result of metapath instances belonging to \mathbf{m} . Otherwise, we use $z_{\mathbf{m}} \times conf_{\mathbf{m} \rightarrow q}$ as the edge-weight, where the metapath \mathbf{m} is mapped to

Algorithm 2 Estimating KG embeddings with Walk-and-Relate (WAR) framework

Input: Graph \mathcal{G} , maximum length of random-walk l_{max} , triplet-loss function f , parameter-sharing strategy s , dictionary with metapaths as key and metapath-information as value \mathcal{D}_1 and dictionary with metapaths as key and rulemaps as value \mathcal{D}_2 .

Output: Embeddings for each node and relation in \mathcal{G}

```

1: Initialize mapping from metapaths to new relations  $N = \{\}$ 
2: for each minibatch  $z$  of nodes in  $\mathcal{G}$  do
3:   Initialize minibatch-triplets  $T = \{\}$ 
4:   for each node  $n_1$  in minibatch  $z$  do
5:     Perform a random walk of length  $l_{max}$  starting from  $n_1$  :
      $[n_1, n_2, \dots, n_{l_{max}}]$ .
6:     for each pair of nodes  $(n_i, n_j)$  in the random walk
      $[n_1, n_2, \dots, n_{l_{max}}]$  such that  $j - i \geq 2$  do
7:       Let the path from  $n_i$  to  $n_j$  describe the metapath  $m$ .
8:       if  $m \in \mathcal{D}_1$ , i.e.,  $m$  is informative with metapath-
       information  $z_m$  then
9:         Get  $rulemap(m)$  of  $m$  from  $\mathcal{D}_2$ 
10:        if  $rulemap(m)$  is non-empty then
11:          Sample a relation  $q$  out of  $rulemap(m)$ , with the
          sampling probability  $conf_{m \rightarrow q}$ .
12:          Add triplet  $(n_i, q, n_j)$  with edge-weight  $z_m \times$ 
           $conf_{m \rightarrow q}$  to the minibatch-triplets set  $T$ .
13:        else
14:          if  $m \in N$  then
15:             $q_{new} = N[m]$ .
16:          else
17:            Introduce a new relation  $q_{new}$  for  $m$ .
18:            Update  $N$  with  $N[m] = q_{new}$ .
19:          end if
20:          Add triplet  $(n_i, q_{new}, n_j)$  with edge-weight  $z_m$  to
          the minibatch-triplets set  $T$ .
21:        end if
22:      end if
23:    end for
24:  end for
25:  Get a sample of edges from the graph  $\mathcal{G}$  and add to  $T$ .
26:  Optimize the given triplet loss function  $f$  and update the
  node and relation embeddings using the triplet minibatch  $T$ 
  and parameter sharing strategy  $s$ .
27: end for
28: return Embeddings for the nodes and relations in  $\mathcal{G}$ 

```

the relation q with rule-confidence $conf_{m \rightarrow q}$. Algorithm 2 provides the pseudo-code for Walk-and-Relate.

5 EXPERIMENTAL METHODOLOGY

We use the Deep Graph Library - Knowledge Embeddings (DGLKE) [41] toolkit to bootstrap the implementation of WAR framework. DGLKE is designed for learning at scale. It provides implementations and benchmarks on widely-used and popular KGE models, making it more easier for users to apply and develop on top of those methods.

Table 2: Knowledge graph datasets.

Dataset	#Nodes	#Edges	#Relations
FB15K	14,951	592,213	1,345
WN18	40,943	151,442	18




Figure 1: Degree distribution of the datasets (training split)

We plan to open-source the code for the WAR framework. We have also open-sourced the code for the WAR framework on GitHub⁴.

5.1 Datasets

We used two standard benchmark datasets (WN18 and FB15K) to evaluate the performance of the proposed approaches and compare them against the baselines. FB15k is derived from the Freebased Knowledge Graph [4], whereas WN18 was derived from WordNet [24]. Table 2 shows various statistics for these datasets. Figure 1 shows the degree distribution of the two datasets. The degree distribution of FB15K is relatively more uniform (has low variance) as compared to WN18. For WN18, $\approx 60\%$ of the nodes have degree ≤ 5 , and $\approx 16\%$ of the nodes have degree ≤ 2 . In contrast, FB15k has $\approx 6\%$ of the nodes having degree ≤ 5 , and $\approx 2\%$ of the nodes having degree ≤ 2 . Both datasets and their training/validation/test splits were used as provided by the DGL-KE toolkit.

5.2 Metrics

We use three standard ranking metrics [21] widely used in literature for evaluating the KG embeddings: Hit@ k (for $k \in \{1, 3, 10\}$), Mean Rank (MR), and Mean Reciprocal Rank (MRR). For a positive triplet i , let S_i be the list of triplets containing i and its associated negative triplets ordered in non-increasing score order, and let $rank_i$ be i th position in S_i . Given that, Hit@ k is the average number of times the positive triplet is among the k highest ranked triplets; MR is the average rank of the positive instances, whereas MRR is the average reciprocal rank of the positive instances. Formally, these metrics are defined as $Hit@k = \frac{1}{Q} \sum_{i=1}^Q \mathbb{1}(rank_i \leq k)$, $MR = \frac{1}{Q} \sum_{i=1}^Q rank_i$, and $MRR = \frac{1}{Q} \sum_{i=1}^Q \frac{1}{rank_i}$, where Q is the total number of positive triplets. Note that Hit@ k and MRR are between 0 and 1, whereas MR ranges from 1 to the $\sum_{i=1}^Q S_i$.

5.3 Comparison methodology

We choose two representative KG embedding approaches based on triplet-loss functions: TransE and Distmult. TransE is a translation-based model and allows us to infer compositional patterns, while Distmult is a similarity scoring-based approach. For these two

⁴<https://github.com/gurdaspuriya/walk-and-relate>

approaches, we evaluate the different triplet-augmentation and parameter-sharing approaches as described below:

- (1) No augmentation: The models are trained on explicit triplets present in the graph, and thus, no augmentation is done.
- (2) WAR - Only rules: To show the promise of proposed augmentation approaches over earlier approaches that only leverage rules of the form $m \rightarrow q$, we provide results on an augmentation exercise where we only add additional triplets based on mapping the metapaths to the existing relations in the graph (to mimic the prior approaches).
- (3) WAR - Metapaths + No parameter sharing (WAR + MP): This approach performs augmentation with the discovered informative-metapaths, by introducing a new relation for each metapath, i.e. without any parameter sharing among the metapath (relation) embeddings.
- (4) WAR - Metapaths + parameter sharing using model’s composition properties (WAR + MP + Model): This approach performs augmentation with the discovered informative-metapaths, and uses the relation embeddings to compute the metapath embeddings based on the underlying composition model. Note that this approach cannot be used for DistMult as it does not allow inference of the composition patterns.
- (5) WAR - Metapaths + sequence modeling for metapath (WAR + MP + RNN): This approach performs augmentation with the discovered informative-metapaths, and uses RNN as the sequence model to aggregate the representations of the constituent relations into the representation of the metapath.
- (6) WAR - Metapaths + basis decomposition (WAR + MP + Basis): This approach performs augmentation with the discovered informative-metapaths, and uses the linear combination of common bases to estimate the embeddings of the metapaths.

5.4 Hyperparameter selection

We use metapaths of lengths up to three ($l_{max} = 3$) to construct augmentation triplets for the WAR approach. All approaches use the mini-batch size of 1,024 nodes to construct random walks. For both WN18 and FB15K, we use the metapath-information threshold of 0.2 to prune the search space. We use the full graph to estimate informative metapaths for WN18 but sample 10% edges for FB15K (given its size). To construct the rules, the rule-confidence threshold is set to 0.5. We use random grid search to optimize for the embedding learning rate (between $[1.0, 1e-1, 1e-2, 1e-3, 1e-4, 1e-5]$), RNN learning rate (between $[1.0, 1e-1, 1e-2, 1e-3, 1e-4, 1e-5]$), basis learning rate (between $[1.0, 1e-1, 1e-2, 1e-3, 1e-4, 1e-5]$) and regularization norm (between $[0, 1e-10, 1e-9, 1e-8, 1e-7, 1e-6]$). We use different learning rates for embeddings, RNN parameters and basis parameters because RNN and basis parameters get dense updates, while embeddings get sparse updates. We use MRR as the primary metric and train the KGE models until the performance on the MRR metric does not increase in two successive epochs. All other hyperparameters are kept the same as provided by DGL-KE for each combination of the triplet loss function and Dataset.

6 RESULTS AND DISCUSSION

Tables 3 and 4 show the performance of various approaches on the WN18 dataset, using the TransE and DistMult scoring functions, respectively. On all the metrics, the approaches that leverage data

augmentation outperform the approaches without data augmentation. Particularly, we observe that leveraging the proposed augmentation framework (WAR) and limiting the augmentation triplets to rules (WAR - Only rules) leads to 0.32% and 0.14% improvement on MRR over the approach without augmentation, for the TransE and DistMult scoring functions, respectively. The improvement increases to 2.92% for the TransE and 0.30% for DistMult when we expand the augmentation triplets with *informative* metapaths, even if those metapaths cannot be mapped to existing relations in the graph, by introducing a new relation for such metapaths. This provides evidence for our hypothesis that we can leverage the composite relations (i.e., metapaths) to provide strong learning signals to the KG embedding approaches, even though we cannot learn confident rules from these metapaths.

Further, given that the number of distinct relations for WN18 is relatively small (18), we do not expect the exponential search space of metapaths to play a role in dampening the performance, thus, we do not expect much performance gain with parameter sharing approaches. This is verified from the experimental results, where we observe that leveraging the parameter sharing is either better or as good as the situation of not performing parameter sharing. With parameter sharing, performance on the WN18 for TransE tends to remain similar, but we see a performance gain of up to 2.92% 1.31% for DistMult.

Tables 5 and 6 show the performance of various approaches on the FB15k dataset, using the TransE and DistMult scoring functions, respectively. Here, we see some different trends as compared with the WN18 dataset. Particularly, we see that using the WAR framework and limiting the augmentation triplets to rules (WAR - Only rules) does not lead to visible performance gain over the approach without augmentation. Given that FB15k is a relatively larger dataset as compared to WN18 (3X lesser nodes and 4X more edges), the limited triplets added by rule-based augmentation do not provide enough signal as compared to the already available triplets in the graph, which explains this behavior. This is also expressed by the degree distribution plots of the two datasets in Figure 1. WN18 has a much larger fraction of nodes with a small degree ($\approx 60\%$ of the nodes have degree ≤ 5) as compared to FB15k ($\approx 6\%$ of the nodes have degree ≤ 5). We see a similar trend when we use augmentation triplets with *informative* metapaths as new relations. This can be attributed to a large number of relations (1,345) in the FB15k dataset, which brings the curse of dimensionality given the exponential search space for metapaths. This is further evident from the performance gain with the parameter-sharing strategies. The experimental results show that leveraging the parameter sharing leads to a performance gain of 1.4350% and 2.7277% on the MRR metric for the TransE and DistMult scoring functions, respectively.

Figure 2 shows the performance metrics for the FB15k dataset and DistMult after sampling the training dataset with different sampling probabilities: 20%, 40%, 60%, 80% and 100% (original, unsampled graph). We observe that the WAR framework provides greater benefits when sampling probabilities are lower for different metrics. The WAR framework’s metrics and those provided by the approach without data augmentation start converging as the sampling probability increases. To facilitate visualization, we also provide relative performance gains based on baseline approaches, i.e., approaches that do not incorporate any data augmentation. These results are




Figure 2: Absolute performance metrics on the FB15k dataset and DistMult model with sampled training triplets.




Figure 3: Percentage performance improvement on the FB15k dataset and DistMult model with sampled training triplets.

provided in figure 3. For example, on the MRR⁵ metric, the WAR-framework-based approaches provide more than 20% performance gain when the training triplets are sampled with a probability of 20%, but the gain drops to 2.7277% for the case when we do not perform sampling. Additionally, limiting the augmentation triplets to rules (WAR - Only rules) leads to minimal gains at 20%, which further diminishes at higher sampling probabilities. This supports our hypothesis that we can leverage metapaths to provide strong learning signals to sparse KG embedding approaches.

The performance metrics are drastically affected by the added sparsity in the training triplets. MRR, for example, scores ≈ 0.8 when trained on unsampled triplets, but this performance drops to < 0.2 when trained with 20% of the training triplets. The drastic decrease is not just because of decreased training data, but because of an increased false-negative rate among the chosen negative samples. KG embedding involves negative sampling to pick out non-observed negative triplets from training data (by corrupting the entities or relations of an actually observed triplet). The chances of selecting an actual positive triplet as a negative triplet (since the positive triplet is absent from the sampled graph) increase as sparsity increases. The effect is even more pronounced with approaches that use self-adversarial negative sampling [29]. By using the KGE model predictions during training, self-adversarial negative sampling selects *hard* negative samples (unobserved samples that the KGE model predicts as positive). Due to the sparsity in the training triplets, the false negative rate increases, not only violating the assumptions of self-adversarial negative sampling but also causing

⁵The MR metric is unbounded, so we observe relatively high variance, and do not observe smooth patterns in the performance gain with the MR, as compared with the other considered metrics.

Table 3: Performance on the WN18 dataset and TransE.

Augmentation	MRR	MR	HITS@1	HITS@3	HITS@10
None	0.5886	211	0.3515	0.8126	0.9435
WAR - Only rules	0.5905	211	0.3548	0.8135	0.9439
WAR + MP	<u>0.6058</u>	186	0.3841	0.8110	0.9413
WAR + MP + Model	0.5970	195	0.3684	0.8082	0.9426
WAR + MP + RNN	<u>0.6047</u>	176	0.3813	0.8115	0.9406
WAR + MP + Basis	<u>0.6041</u>	190	0.3800	0.8120	0.9415

Table 4: Performance on the WN18 dataset and DistMult.

Augmentation	MRR	MR	HITS@1	HITS@3	HITS@10
None	0.8290	655	0.7351	0.9204	0.9451
WAR - Only rules	0.8302	665	0.7372	0.9197	0.9457
WAR + MP	0.8315	555	0.7385	0.9205	0.9478
WAR + MP + RNN	<u>0.8399</u>	571	0.7611	0.9151	0.9407
WAR + MP + Basis	<u>0.8353</u>	592	0.7446	0.9224	0.9479

Table 5: Performance on the FB15k dataset and TransE.

Augmentation	MRR	MR	HITS@1	HITS@3	HITS@10
None	0.6899	37	0.5810	0.7734	0.8603
WAR - Only rules	0.6899	37	0.5810	0.7737	0.8603
WAR + MP	0.6843	39	0.5746	0.7680	0.8585
WAR + MP + Model	0.6937	37	0.5803	0.7820	0.8735
WAR + MP + RNN	<u>0.6998</u>	36	0.5881	0.7869	0.8766
WAR + MP + Basis	<u>0.6967</u>	37	0.5847	0.7836	0.8753

positive triplets to be chosen as negative triplets more frequently than random triplets, which will result in poor performance.

7 CONCLUSION

In this paper, we address the challenge of data sparsity that usually occurs in real-world knowledge graphs, i.e., the lack of enough

Table 6: Performance on the FB15k dataset and DistMult.

Augmentation	MRR	MR	HITS@1	HITS@3	HITS@10
None	0.7992	117	0.7595	0.8240	0.8703
WAR - Only rules	0.7974	96	0.7546	0.8237	0.8748
WAR + MP	<u>0.8053</u>	94	0.7648	0.8309	0.8761
WAR + MP + RNN	<u>0.8107</u>	93	0.7717	0.8368	0.8765
WAR + MP + Basis	<u>0.8210</u>	77	0.7876	0.8413	0.8795

triplets per entity. Specifically, we propose an efficient data augmentation approach, named Walk-and-Relate (WAR), to augment the number of triplets. WAR leverages random walks to create additional triplets, such that the relations carried by these introduced triplets entail the metapath induced by the random walks. We provide approaches to accurately and efficiently filter out informative metapaths from the possible set of metapaths, induced by the random walks. The proposed approaches are model-agnostic, and the augmented data can be used with any KG embedding approach out of the box. Experimental results obtained on the benchmark datasets show the advantages of the proposed approach.

Our work makes a step towards going beyond inference rules to augment KGs, and envision that the proposed method will serve as a motivation to develop other solutions to address the sparsity in real-world KGs.

REFERENCES

- [1] Bilal Abu-Salih, Marwan Al-Tawil, Ibrahim Aljarah, Hossam Faris, Pornpit Wongthongtham, Kit Yan Chan, and Amin Beheshti. 2021. Relational learning analysis of social politics using knowledge graph embedding. *Data Mining and Knowledge Discovery* 35, 4 (2021), 1497–1536.
- [2] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. 1993. Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*. 207–216.
- [3] Molood Barati, Quan Bai, and Qing Liu. 2016. SWARM: an approach for mining semantic association rules from semantic web data. In *Pacific rim international conference on artificial intelligence*. Springer, 30–43.
- [4] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. 1247–1250.
- [5] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems* 26 (2013).
- [6] Richard P. Brent. 1971. An algorithm with guaranteed convergence for finding a zero of a function. *The computer journal* 14, 4 (1971), 422–425.
- [7] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [8] Michael Fell, Elena Cabrio, Maroua Tikat, Franck Michel, Michel Buffa, and Fabien Gandon. 2022. The WASABI song corpus and knowledge graph for music lyrics analysis. *Language Resources and Evaluation* (2022), 1–31.
- [9] Luis Galárraga, Christina Teflioudi, Katja Hose, and Fabian M Suchanek. 2015. Fast rule mining in ontological knowledge bases with AMIE++. *The VLDB Journal* 24, 6 (2015), 707–730.
- [10] Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian Suchanek. 2013. AMIE: association rule mining under incomplete evidence in ontological knowledge bases. In *Proceedings of the 22nd international conference on World Wide Web*. 413–422.
- [11] Palash Goyal and Emilio Ferrara. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* 151 (2018), 78–94.
- [12] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [13] Shu Guo, Quan Wang, Lihong Wang, Bin Wang, and Li Guo. 2016. Jointly embedding knowledge graphs and logical rules. In *Proceedings of the 2016 conference on empirical methods in natural language processing*. 192–202.
- [14] Vinh Thinh Ho, Daria Stepanova, Mohamed H Gad-Elrab, Evgeny Kharlamov, and Gerhard Weikum. 2018. Rule learning from knowledge graphs guided by embedding models. In *International Semantic Web Conference*. Springer, 72–90.
- [15] Vassilis N Ioannidis, Xiang Song, Saurav Manchanda, Mufei Li, Xiaoqin Pan, Da Zheng, Xia Ning, Xiangxiang Zeng, and George Karypis. 2020. Drkg-drug repurposing knowledge graph for covid-19. *arXiv preprint arXiv: 2010.09600* (2020).
- [16] Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd annual meeting of the association for computational linguistics and the 7th international joint conference on natural language processing (volume 1: Long papers)*. 687–696.
- [17] Jack Kiefer. 1953. Sequential minimax search for a maximum. *Proceedings of the American mathematical society* 4, 3 (1953), 502–506.
- [18] Denis Krompaß, Stephan Baier, and Volker Tresp. 2015. Type-constrained representation learning in knowledge graphs. In *International semantic web conference*. Springer, 640–655.
- [19] Jonathan Lajus, Luis Galárraga, and Fabian Suchanek. 2020. Fast and exact rule mining with AMIE 3. In *European Semantic Web Conference*. Springer, 36–52.
- [20] Ni Lao, Tom Mitchell, and William Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *Proceedings of the 2011 conference on empirical methods in natural language processing*. 529–539.
- [21] Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. Pytorch-biggraph: A large scale graph embedding system. *Proceedings of Machine Learning and Systems* 1 (2019), 120–131.
- [22] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*.
- [23] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems* 26 (2013).
- [24] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.
- [25] Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2017. A novel embedding model for knowledge base completion based on convolutional neural network. *arXiv preprint arXiv:1712.02121* (2017).
- [26] Fabrizio Orlandi, Jeremy Debattista, Islam A Hassan, Clare Conran, Majid Latifi, Matthew Nicholson, Fahim A Salim, Daniel Turner, Owen Conlan, Declan O’sullivan, et al. 2018. Leveraging knowledge graphs of movies and their content for web-scale analysis. In *2018 14th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*. IEEE, 609–616.
- [27] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [28] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European semantic web conference*. Springer, 593–607.
- [29] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197* (2019).
- [30] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *International conference on machine learning*. PMLR, 2071–2080.
- [31] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (2017), 2724–2743.
- [32] Zhichun Wang and Juanzi Li. 2015. RDF2Rules: Learning rules from RDF knowledge bases by mining frequent predicate cycles. *arXiv preprint arXiv:1512.07734* (2015).
- [33] Zhigang Wang, Juanzi Li, Zhiyuan Liu, and Jie Tang. 2016. Text-enhanced representation learning for knowledge graph. In *Proceedings of International Joint Conference on Artificial Intelligent (IJCAI)*. 4–17.
- [34] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 28.
- [35] Ruobing Xie, Zhiyuan Liu, Maosong Sun, et al. 2016. Representation Learning of Knowledge Graphs with Hierarchical Types.. In *IJCAI*, Vol. 2016. 2965–2971.
- [36] Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. DeepPath: A reinforcement learning method for knowledge graph reasoning. *arXiv preprint arXiv:1707.06690* (2017).
- [37] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575* (2014).
- [38] Nasser Zalmout, Chenwei Zhang, Xian Li, Yan Liang, and Xin Luna Dong. 2021. All You Need to Know to Build a Product Knowledge Graph. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 4090–4091.
- [39] Wen Zhang, Bibek Paudel, Liang Wang, Jiaoyan Chen, Hai Zhu, Wei Zhang, Abraham Bernstein, and Huajun Chen. 2019. Iteratively learning embeddings and rules for knowledge graph reasoning. In *The World Wide Web Conference*.

2366–2377.

- [40] Feng Zhao, Haoran Sun, Langjunqing Jin, and Hai Jin. 2020. Structure-augmented knowledge graph embedding for sparse data with rule learning. *Computer Communications* 159 (2020), 271–278.
- [41] Da Zheng, Xiang Song, Chao Ma, Zeyuan Tan, Zihao Ye, Jin Dong, Hao Xiong, Zheng Zhang, and George Karypis. 2020. Dgl-ke: Training knowledge graph embeddings at scale. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 739–748.