

# Knowledge graph embedding for data mining vs. knowledge graph embedding for link prediction – two sides of the same coin?

Jan Portisch<sup>a,b</sup>, Nicolas Heist<sup>a</sup> and Heiko Paulheim<sup>a,\*</sup>

<sup>a</sup> *Data and Web Science Group, University of Mannheim, Germany*

*E-mails: [jan@informatik.uni-mannheim.de](mailto:jan@informatik.uni-mannheim.de), [nico@informatik.uni-mannheim.de](mailto:nico@informatik.uni-mannheim.de),*

*[heiko@informatik.uni-mannheim.de](mailto:heiko@informatik.uni-mannheim.de)*

<sup>b</sup> *SAP SE, Germany*

*E-mail: [jan.portisch@sap.com](mailto:jan.portisch@sap.com)*

**Editors:** Mehwish Alam, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Germany; Davide Buscaldi, LIPN, Université Sorbonne Paris Nord, France; Michael Cochez, Vrije University of Amsterdam, the Netherlands; Francesco Osborne, Knowledge Media Institute, (KM<sub>i</sub>), and The Open University, UK; Diego Reforgiato Recupero, University of Cagliari, Italy; Harald Sack, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Germany

**Solicited reviews:** Angelo Salatino, The Open University, UK; Paul Groth, University of Amsterdam, Netherlands; Five anonymous reviewers

**Abstract.** Knowledge Graph Embeddings, i.e., projections of entities and relations to lower dimensional spaces, have been proposed for two purposes: (1) providing an encoding for data mining tasks, and (2) predicting links in a knowledge graph. Both lines of research have been pursued rather in isolation from each other so far, each with their own benchmarks and evaluation methodologies. In this paper, we argue that both tasks are actually related, and we show that the first family of approaches can also be used for the second task and vice versa. In two series of experiments, we provide a comparison of both families of approaches on both tasks, which, to the best of our knowledge, has not been done so far. Furthermore, we discuss the differences in the similarity functions evoked by the different embedding approaches.

**Keywords:** Knowledge graph embedding, link prediction, data mining, RDF2vec, KGlove, node2vec, DeepWalk, TransE, TransR, RotatE, RESCAL, DistMult, ComplEx

## 1. Introduction

In the recent past, the topic of knowledge graph embedding – i.e., projecting entities and relations in a knowledge graph into a numerical vector space – has gained a lot of traction. An often cited survey from 2017 [42] lists already 25 approaches, with new models being proposed almost every month, as depicted in Fig. 1.

Even more remarkably, two mostly disjoint strands of research have emerged in that vivid area. The first family of research works focus mostly on *link prediction* [12], i.e., the approaches are evaluated in a knowledge graph refinement setting [24]. The optimization goal here is to distinguish correct from incorrect triples in the knowledge graph

---

\* Corresponding author. E-mail: [heiko@informatik.uni-mannheim.de](mailto:heiko@informatik.uni-mannheim.de).

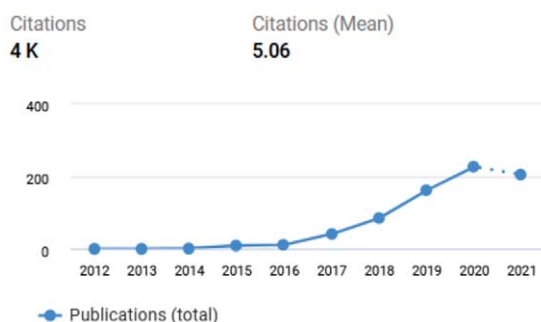


Fig. 1. Publications with *Knowledge Graph Embedding* in their title or abstract, created with dimensions.ai (as of November 15th, 2021).

as accurately as possible. The evaluations of this kind of approaches are always conducted within the knowledge graph, using the existing knowledge graph assertions as ground truth.

A second strand of research focuses on the embedding of entities in the knowledge graph for downstream tasks outside the knowledge graph, which often come from the data mining field – hence, we coin this family of approaches *embeddings for data mining*. Examples include: the prediction of external variables for entities in a knowledge graph [30], information retrieval backed by a knowledge graph [36], or the usage of a knowledge graph in content-based recommender systems [32]. In those cases, the optimization goal is to create an embedding space which reflects *semantic similarity* as well as possible (e.g., in a recommender system, similar items to the ones in the user interest should be recommended). The evaluations here are always conducted outside the knowledge graph, based on external ground truth.

In this paper, we want to look at the commonalities and differences of the two approaches. We look at two of the most basic and well-known approaches of both strands, i.e., *TransE* [4] and *RDF2vec* [30], and analyze and compare their optimization goals in a simple example. Moreover, we analyze the performance of approaches from both families in the respective other evaluation setup: we explore the usage of link-prediction based embeddings for other downstream tasks based on similarity, and we propose a link prediction method based on node embedding techniques such as *RDF2vec*. From those experiments, we derive a set of insights into the differences of the two families of methods, and a few recommendations on which kind of approach should be used in which setting.

## 2. Related work

As pointed out above, the number of works on knowledge graph embedding is legion, and enumerating them all in this section would go beyond the scope of this paper. However, there have already been quite a few survey articles.

The first strand of research works – i.e., knowledge graph embeddings for link prediction – has been covered in different surveys, such as [42], and, more recently, [8,14,33]. The categorization of approaches in those reviews is similar, as they distinguish different families of approaches: *translational distance models* [42] or *geometric models* [33] focus on link prediction as a geometric task, i.e., projecting the graph in a vector space so that a translation operation defined for relation  $r$  on a head  $h$  yields a result close to the tail  $t$ .

The second family among the link prediction embeddings are *semantic matching* [42] or *matrix factorization or tensor decomposition* [33] models. Here, a knowledge graph is represented as a three-dimensional tensor, which is decomposed into smaller tensors and/or two-dimensional matrices. The reconstruction operation can then be used for link prediction.

The third and youngest family among the link prediction embeddings are based on deep learning and graph neural networks. Here, neural network training approaches, such as convolutional neural networks, capsule networks, or recurrent neural networks, are adapted to work with knowledge graphs. They are generated by training a deep neural network. Different architectures exist (based on convolutions, recurrent layers, etc.), and the approaches also differ in the training objective, e.g., performing binary classification into true and false triples, or predicting the relation of a triple, given its subject and object. [33].

While most of those approaches only consider graphs with nodes and edges, most knowledge graphs also contain literals, e.g., strings and numeric values. Recently, approaches combining textual information with knowledge graph embeddings using language modeling techniques have also been proposed, using techniques such as word2vec and convolutional neural networks [45] or transformer methods [9,43]. [11] shows a survey of approaches which take such literal information into account. It is also one of the few review articles which considers embedding methods from the different research strands.

Link prediction is typically evaluated on a set of standard datasets, and uses a within-KG protocol, where the triples in the knowledge graph are divided into a training, testing, and validation set. Prediction accuracy is then assessed on the validation set. Datasets commonly used for the evaluation are FB15k, which is a subset of Freebase, and WN18, which is derived from WordNet [4]. Since it has been remarked that those datasets contain too many simple inferences due to inverse relations, the more challenging variants FB15k-237 [39] and WN18RR [10] have been proposed. More recently, evaluation sets based on larger knowledge graphs, such as YAGO3-10 [10] and DBpedia50k/DBpedia500k [34] have been introduced.

The second strand of research works, focusing on the embedding for downstream tasks (which are often from the domain of data mining), is not as extensively reviewed, and the number of works in this area are still smaller. One of the more comprehensive evaluations is shown in [7], which is also one of the rare works which includes approaches from both strands in a common evaluation. They show that at least the three methods for link prediction used – namely TransE, TransR, and TransH – perform inferior on downstream tasks, compared to approaches developed specifically for optimizing for entity similarity in the embedding space.

A third, yet less closely related strand of research works is *node embeddings for homogeneous graphs*, such as node2vec and DeepWalk. While knowledge graphs come with different relations and are thus considered *heterogeneous*, approaches for homogeneous graphs are sometimes used on knowledge graphs as well by first transforming the knowledge graph into an unlabeled graph, usually by ignoring the different types of relations. Since some of the approaches are defined for undirected graphs, but knowledge graphs are directed, those approaches may also ignore the direction of edges.

For the evaluation of entity embeddings for data mining, i.e., optimized for capturing entity similarity, there are quite a few use cases at hand. The authors in [25] list a number of tasks, including classification and regression of entities based on external ground truth variables, entity clustering, as well as identifying semantically related entities.

Most of the above mentioned strands exist mainly in their own respective “research bubbles”. Table 1 shows a co-citation analysis of the different families of approaches. It shows that the Trans\* family, together with other approaches for link prediction, forms its own citation network, so do the approaches for homogeneous networks, while RDF2vec and KGlove are less clearly separated.

Works which explicitly compare approaches from the different research strands are still rare. In [48], the authors analyze the vector spaces of different embedding models with respect to class separation, i.e., they fit the best linear

Table 1

Co-citation likelihood of different embeddings approaches, obtained from Google scholar, July 12th, 2021. An entry (*row,column*) in the table reads as: *this fraction of the papers citing column also cites row*

	TransE	TransR	RotatE	DistMult	RESCAL	Complex	RDF2vec	KGlove	node2vec	DeepWalk
<i>total</i>	3379	1852	391	1147	408	1017	321	73	5269	5290
<b>TransE</b>	100,0%	69,7%	80,6%	69,8%	40,9%	73,3%	38,3%	46,6%	5,1%	5,0%
<b>TransR</b>	32,6%	100,0%	36,3%	36,2%	23,8%	38,2%	17,8%	27,4%	3,2%	2,9%
<b>RotatE</b>	10,7%	10,7%	100,0%	22,8%	3,4%	27,3%	3,1%	6,8%	0,5%	0,5%
<b>DistMult</b>	25,2%	26,7%	75,4%	100,0%	16,7%	63,7%	16,2%	21,9%	1,7%	1,5%
<b>RESCAL</b>	22,3%	27,4%	34,0%	38,6%	100,0%	40,8%	14,3%	20,5%	1,6%	1,7%
<b>Complex</b>	26,4%	27,2%	73,4%	58,2%	19,9%	100,0%	15,3%	17,8%	1,9%	1,9%
<b>RDF2vec</b>	4,4%	4,6%	4,3%	4,8%	2,5%	5,9%	100,0%	83,6%	1,6%	1,5%
<b>KGlove</b>	1,0%	1,4%	2,0%	1,0%	1,2%	1,4%	13,7%	100,0%	0,3%	0,2%
<b>node2vec</b>	11,1%	11,7%	8,7%	9,4%	5,6%	9,7%	21,5%	31,5%	100,0%	64,1%
<b>DeepWalk</b>	11,7%	11,7%	10,0%	8,8%	6,1%	9,1%	21,8%	23,3%	66,2%	100,0%

separation between classes in different embedding spaces. According to their findings, RDF2vec achieves a better linear separation than the models tailored to link prediction.

In [6], an in-KG scenario, i.e., the detection and correction of erroneous links, is considered. The authors compare RDF2vec (with an additional classification layer) to TransE and DistMult on the link prediction task. The results are mixed: While RDF2vec outperforms TransE and DistMult in terms of Mean Reciprocal Rank and Precision@1, it is inferior in Precision@10. Since the results are only validated on one single dataset, the evidence is rather thin.

Most other research works in which approaches from different strands are compared are related to different downstream tasks. In many cases, the results are rather inconclusive, as the following examples illustrate:

- [5] and [15] both analyze *drug drug interaction*, using different sets of embedding methods. The finding of [5] is that RDF2vec outperforms TransE and TransD, whereas in the experiment in [15], ComplEx outperforms RDF2vec, KGlove, TransE, and CrossE, and, in particular, TransE outperforms RDF2vec.
- [3,6], and [44] all analyze *link prediction in different graphs*. While [3] state that RotatE and TransD outperform TransE, DistMult, and ComplEx, which in turn outperforms node2vec, [6] reports that DistMult outperforms RDF2vec, which in turn outperforms TransE, while [44] reports that KG2vec (which can be considered equivalent to RDF2vec) outperforms node2vec, which in turn outperforms TransE.
- [2] compare the performance of RDF2vec, DistMult, TransE, and SimpleE on a set of classification and clustering datasets. The results are mixed. For classification, the authors use four different learning algorithms, and the variance induced by the learning algorithms is most often higher than that induced by the embedding method. For the clustering, they report that TransE outperforms the other approaches.<sup>1</sup>

While this is not a comprehensive list, these observations hint at a need both for more task-specific benchmark datasets as well as for ablation studies analyzing the interplay of embedding methods and other processing steps. Moreover, it is important to gain a deeper understanding of how these approaches behave with respect to different downstream problems, and to have more direct comparisons. This paper aims at closing the latter gap.

### 3. Knowledge graph embedding methods for data mining

Traditionally, most data mining methods are working on propositional data, i.e., each instance is a row in a table, described by a set of (binary, numeric, or categorical) features. For using knowledge graphs in data mining, one needs to either develop methods which work on graphs instead of propositional data, or find ways to represent instances of the knowledge graph as feature vectors [31]. The latter is often referred to as *propositionalization* [29].

RDF2vec [30] is a prominent example from the second family. It adapts the word2vec approach [22] for deriving word embeddings (i.e., vector representations for words) from a corpus of sentences. RDF2vec creates such sentences by performing *random walks* on an RDF graph and collecting the sequences of entities and relations, then trains a word2vec model on those sequences. It has been shown that this strategy outperforms other strategies of propositionalization. The relation between propositionalization and embedding methods has also recently been pointed out by [17].

#### 3.1. Data mining is based on similarity

Predictive data mining tasks are predicting classes or numerical values for instances. Typically, the target is to predict an external variable not contained in the knowledge graph (or, to put it differently: use the background information from the knowledge graph to improve prediction models). One example would be to predict the popularity of an item (e.g., a book, a music album, a movie) as a numerical value. The idea here would be that two items which share similar features should also receive similar ratings. The same mechanism is also exploited in recommender

---

<sup>1</sup>We think that these results must be taken with a grain of salt. To evaluate the clustering quality, the authors use an intrinsic evaluation metric, i.e., Silhouette score, which is computed in the respective vector space. It is debatable, however, whether Silhouette scores computed in different vector spaces are comparable.

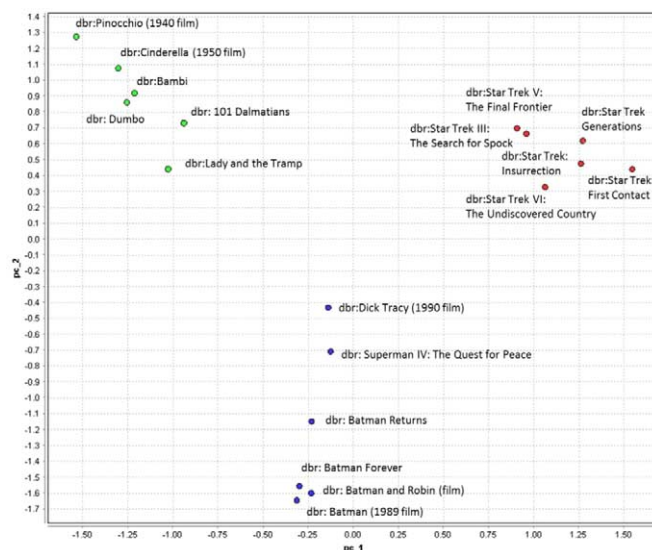


Fig. 2. RDF2vec embeddings for movies in DBpedia, from [32].

systems: if two items share similar features, users who consumed one of those items are recommended the other one.

RDF2vec has been shown to be usable for such cases, since the underlying method tends to create similar vectors for similar entities, i.e., position them closer in vector space [32]. Figure 2 illustrates this using a 2D PCA plot of RDF2vec vectors for movies in DBpedia. It can be seen that clusters of movies, e.g., Disney movies, Star Trek movies, and Marvel related movies are formed.

Many techniques for predictive data mining rely on similarity in one or the other way. This is more obvious for, e.g.,  $k$ -nearest neighbors, where the predicted label for an instance is the majority or average of labels of its closest neighbors (i.e., most similar instances), or Naive Bayes, where an instance is predicted to belong to a class if its feature values are most similar to the typical distribution of features for this class (i.e., it is similar to an average member of this class). A similar argument can be made for neural networks, where one can assume a similar output when changing the value of one input neuron (i.e., one feature value) by a small delta. Other classes of approaches (such as Support Vector Machines) use the concept of *class separability*, which is similar to exploiting similarity: datasets with well separable classes have similar instances (belonging to the same class) close to each other, while dissimilar instances (belonging to different classes) are further away from each other [38].

### 3.2. Creating similar embeddings for similar instances

To understand how (and why) RDF2vec creates embeddings that project similar entities to nearby vectors, we use the running example depicted in Fig. 3 and Fig. 4, showing a number of European cities, countries, and heads of those governments.

As discussed above, the first step of RDF2vec is to create random walks on the graph. To that end, RDF2vec starts a fixed number of random walks of a fixed maximum length from each entity. Since the example above is very small, we will, for the sake of illustration, enumerate *all* walks of length 4 that can be created for the graph. Those walks are depicted in Fig. 5. It is notable that, since the graph has nodes without outgoing edges, some of the walks are actually shorter than 4.

In the next step, the walks are used to train a predictive model. Since RDF2vec uses word2vec, it can be trained with the two flavors of word2vec, i.e., CBOW (context back of words) and SG (skip gram). The first predicts a word, given its surrounding words, the second predicts the surroundings, given a word. For the sake of our argument, we will only consider the second variant, depicted in Fig. 6. Simply speaking, given training examples where the input is the target word (as a one-hot-encoded vector) and the output is the context words (again, one hot encoded vectors),

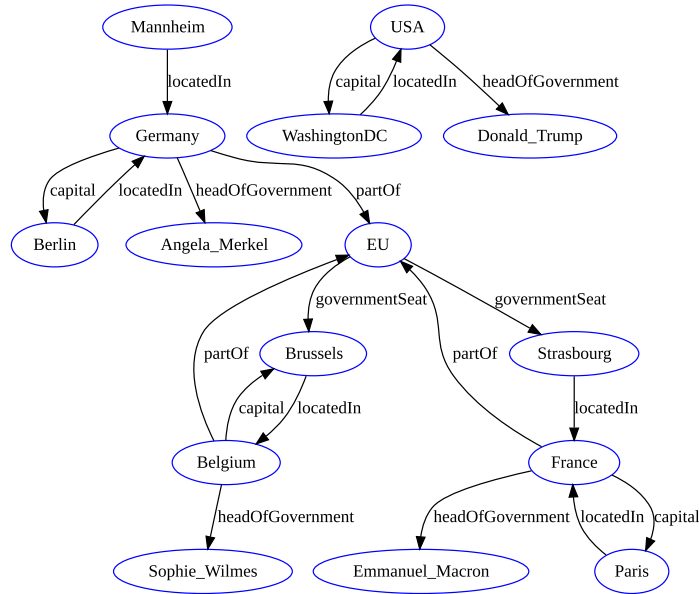


Fig. 3. Example graph used for illustration.

```

Berlin locatedIn Germany .
Germany headOfGovernment Angela_Merkel .
Mannheim locatedIn Germany .
Belgium capital Brussels .
Germany partOf EU .
Belgium partOf EU .
Belgium headOfGovernment Sophie_Wilmes .
EU governmentSeat Brussels .
USA capital WashingtonDC .
WashingtonDC locatedIn USA .
France capital Paris .
France headOfGovernment Emmanuel_Macron .
Paris locatedIn France .
Strasbourg locatedIn France .
Germany capital Berlin .
Brussels locatedIn Belgium .
France partOf EU .
USA headOfGovernment Donald_Trump .
EU governmentSeat Strasbourg .

```

Fig. 4. Triples of the example knowledge graph.

a neural network is trained, where the hidden layer is typically of smaller dimensionality than the input. That hidden layer is later used to produce the actual embedding vectors.

To create the training examples, a window with a given size is slid over the input sentences. Here, we use a window of size 2, which means that the two words preceding and the two words succeeding a context word are taken into consideration. Table 2 shows the training examples generated for three instances.

A model that learns to predict the context given the target word would now learn to predict the *majority* of the context words for the target word at hand at the output layer called *output* in Fig. 6, as depicted in the lower part of Table 2. Here, we can see that Paris and Berlin share two out of four predictions, so do Mannheim and Berlin. Angela Merkel and Berlin share one out of four predictions.<sup>2</sup>

<sup>2</sup>Note that in the classic formulation of RDF2vec (and word2vec), the position at which a prediction appears does not matter. The order-aware variant *RDF2vec<sub>oa</sub>* [27] uses an order-aware formulation of word2vec [20].

```

Belgium partOf EU governmentSeat Brussels
Belgium capital Brussels locatedIn Belgium
Belgium partOf EU governmentSeat Strasbourg
Belgium headOfGovernment Sophie_Wilmes
Berlin locatedIn Germany capital Berlin
Berlin locatedIn Germany headOfGovernment Angela_Merkel
Berlin locatedIn Germany partOf EU
Brussels locatedIn Belgium headOfGovernment Sophie_Wilmes
Brussels locatedIn Belgium partOf EU
Brussels locatedIn Belgium capital Brussels
EU governmentSeat Strasbourg locatedIn France
EU governmentSeat Brussels locatedIn Belgium
France headOfGovernment Emmanuel_Macron
France capital Paris locatedIn France
France partOf EU governmentSeat Brussels
France partOf EU governmentSeat Strasbourg
Germany partOf EU governmentSeat Brussels
Germany partOf EU governmentSeat Strasbourg
Germany capital Berlin locatedIn Germany
Germany headOfGovernment Angela_Merkel
Mannheim locatedIn Germany capital Berlin
Mannheim locatedIn Germany headOfGovernment Angela_Merkel
Mannheim locatedIn Germany partOf EU
Paris locatedIn France headOfGovernment Emmanuel_Macron
Paris locatedIn France partOf EU
Paris locatedIn France capital Paris
Strasbourg locatedIn France capital Paris
Strasbourg locatedIn France headOfGovernment Emmanuel_Macron
Strasbourg locatedIn France partOf EU
USA headOfGovernment Donald_Trump
USA capital Washington_DC locatedIn USA
Washington_DC locatedIn USA capital Washington_DC
Washington_DC locatedIn USA headOfGovernment Donald_Trump
    
```

Fig. 5. Walks extracted from the example graph.

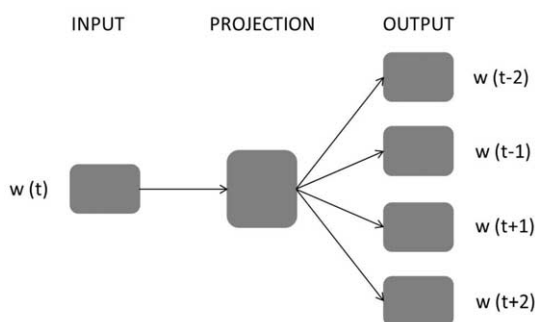


Fig. 6. The skip gram variant of word2vec [30].

Considering again Fig. 6, given that the activation function which computes the *output* from the *projection* values is continuous, it implies that similar activations on the output layer requires similar values on the projection layer. Hence, for a well fit model, the distance on the projection layer of Paris, Berlin, and Mannheim should be comparatively lower than the distance of the other entities, since they activate similar outputs.<sup>3</sup>

Figure 7 depicts a two-dimensional RDF2vec embedding learned for the example graph.<sup>4</sup> We can observe that there are clusters of persons, countries, and cities. The grouping of similar objects also goes further – we can, e.g.,

<sup>3</sup>Note that there are still weights learned for the individual connections between the projection and the output layer, which emphasize some connections more strongly than others. Hence, we cannot simplify our argumentation in a way like “with two common context words activated, the entities must be projected twice as close as those with one common context word activated”.

<sup>4</sup>Created with PyRDF2vec [41], using two dimensions, a walk length of 8, and standard configuration otherwise.

Table 2

Training examples for instances *Paris*, *Berlin*, *Mannheim*, *Angela Merkel*, *Donald Trump*, and *Belgium* (upper part) and majority predictions (lower part)

Target word	$w_{-2}$	$w_{-1}$	$w_{+1}$	$w_{+2}$
Paris	France	capital	locatedIn	France
Paris	–	–	locatedIn	France
Paris	–	–	locatedIn	France
Paris	–	–	locatedIn	France
Paris	France	capital	–	–
Paris	France	capital	–	–
Berlin	–	–	locatedIn	Germany
Berlin	Germany	capital	–	–
Berlin	–	–	locatedIn	Germany
Berlin	–	–	locatedIn	Germany
Berlin	Germany	capital	locatedIn	Germany
Berlin	Germany	capital	–	–
Mannheim	–	–	locatedIn	Germany
Mannheim	–	–	locatedIn	Germany
Mannheim	–	–	locatedIn	Germany
Angela Merkel	Germany	headOfGovernment	–	–
Angela Merkel	Germany	headOfGovernment	–	–
Angela Merkel	Germany	headOfGovernment	–	–
Donald Trump	USA	headOfGovernment	–	–
Donald Trump	USA	headOfGovernment	–	–
Belgium	–	–	partOf	EU
Belgium	–	–	capital	Brussels
Belgium	Brussels	locatedIn	–	–
Belgium	–	–	partOf	EU
Belgium	–	–	headOfGovernment	Sophie Wilmes
Belgium	Brussels	locatedIn	headOfGovernment	Sophie Wilmes
Belgium	Brussels	locatedIn	partOf	EU
Belgium	Brussels	locatedIn	capital	Brussels
Belgium	Brussels	locatedIn	–	–
Paris	France	capital	locatedIn	France
Berlin	Germany	capital	locatedIn	Germany
Mannheim	–	–	locatedIn	Germany
Angela Merkel	Germany	headOfGovernment	–	–
Donald Trump	USA	headOfGovernment	–	–
Belgium	Brussels	locatedIn	partOf	EU

observe that European cities in the dataset are embedded closer to each other than to Washington D.C. This is in line with previous observations showing that RDF2vec is particularly well suited in creating clusters also for finer-grained classes [35]. A predictive model could now exploit those similarities, e.g., for type prediction, as proposed in [16] and [35].

### 3.3. Usage for link prediction

From Fig. 7, we can assume that link prediction should, in principle, be possible. For example, the predictions for heads of governments all point in a similar direction. This is in line with what is known about word2vec, which



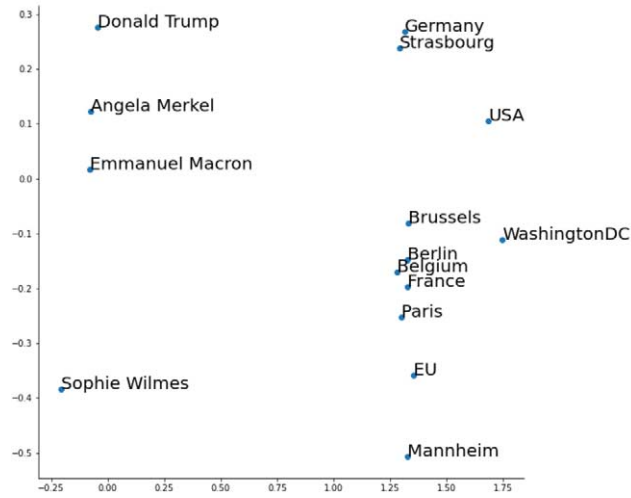


Fig. 7. The example graph embedded with RDF2vec.

allows for computing analogies, like the well-known example

$$v(King) - v(Man) + v(Woman) \approx v(Queen) \tag{1}$$

RDF2vec does not learn relation embeddings, only entity embeddings.<sup>5</sup> Hence, we cannot directly predict links, but we can exploit those analogies. If we want to make a tail prediction like

$$\langle h, r, ? \rangle, \tag{2}$$

we can identify another pair  $\langle h', r, t' \rangle$  and exploit the above analogy, i.e.,

$$t' - h' + h \approx t \tag{3}$$

To come to a stable prediction, we would use the average, i.e.,

$$t \approx \frac{\sum_{\langle h', r, t' \rangle} t' - h' + h}{|\langle h', r, t' \rangle|}, \tag{4}$$

where  $|\langle h', r, t' \rangle|$  is the number of triples which have  $r$  as predicate.

With the same idea, we can also average the relation vectors  $r$  for each relation that holds between all its head and tail pairs, i.e.,

$$r \approx \frac{\sum_{\langle h', r, t' \rangle} t' - h'}{|\langle h', r, t' \rangle|}, \tag{5}$$

and thereby reformulate the above equation to

$$t \approx h + r, \tag{6}$$

which is what we expect from an embedding model for link prediction. Those approximate relation vectors for the example at hand are depicted in Fig. 8. We can see that in some (not all) cases, the directions of the vectors are

<sup>5</sup>Technically, we can also make RDF2vec learn embeddings for the relations, but they would not behave the way we need them.

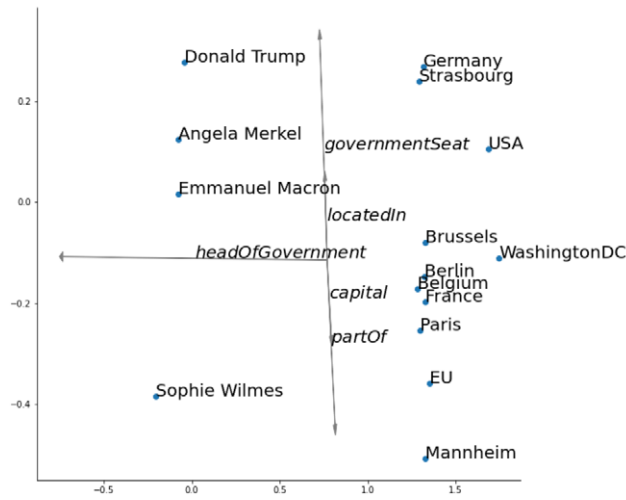


Fig. 8. Average relation vectors for the example.

approximately correct: the *partOf* vector is roughly the difference between *EU* and *Germany*, *France*, and *Belgium*, and the *headOfGovernment* vector is approximately the vector between the countries and the politicians cluster.

It can also be observed that the vectors for *locatedIn* and *capitalOf* point in reverse directions, which makes sense because they form connections between two clusters (countries and cities) in opposite directions.

#### 4. Knowledge graph embedding methods for link prediction

A larger body of work has been devoted on knowledge graph embedding methods for link prediction. Here, the goal is to learn a model which embeds entities and relations in the same vector space.

##### 4.1. Link prediction is based on vector operations

As the main objective is link prediction, most models, more or less, try to find a vector space embedding of entities and relations so that

$$t \approx h \oplus r \quad (7)$$

holds for as many triples  $\langle h, r, t \rangle$  as possible.  $\oplus$  can stand for different operations in the vector space; in basic approaches, simple vector addition (+) is used. In our considerations below, we will also use vector addition.

In most approaches, negative examples are created by corrupting an existing triple, i.e., replace the head or tail with another entity from the graph (some approaches also foresee corrupting the relation). Then, a model is learned which tries to tell apart corrupted from non-corrupted triples. The formulation in the original TransE paper [4] defines the loss function  $L$  as follows:

$$L = \sum_{\substack{(h,r,t) \in S, \\ (h',r,t') \in S'}} [\gamma + d(h+r, t) - d(h'+r, t')]_+ \quad (8)$$

where  $\gamma$  is some margin, and  $d$  is a distance function, i.e., the  $L1$  or  $L2$  norm.  $S$  is the set of statements that are in the knowledge graph, and  $S'$  are the corrupted statements derived from them. In words, the formula states for a triple  $\langle h, r, t \rangle$ ,  $h+r$  should be closer to  $t$  than to  $t'$  for some corrupted tail, similarly for a corrupted head. However, a difference of  $\gamma$  is accepted.

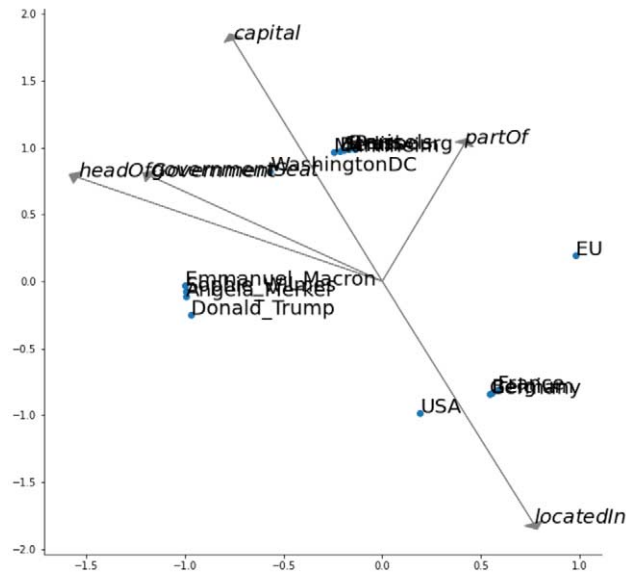


Fig. 9. Example graph embedded by TransE.

Figure 9 shows the example graph from above, as embedded by TransE.<sup>6</sup> Looking at the relation vectors, it can be observed that they seem approximately accurate in some cases, e.g.,

$$Germany + headOfGovernment \approx Angela\_Merkel,$$

but not everywhere.<sup>7</sup>

Like in the RDF2vec example above, we can observe that the two vectors for *locatedIn* and *capital* point in opposite directions. Also similar to the RDF2vec example, we can see that entities in similar classes form clusters: cities are mostly in the upper part of the space, people in the left, countries in the lower right part.

#### 4.2. Usage for data mining

As discussed above, positioning similar entities close in a vector space is an essential requirement for using entity embeddings in data mining tasks. To understand why an approach tailored towards link prediction can also, to a certain extent, cluster similar instances together (although not explicitly designed for this task), we first rephrase the approximate link prediction Equation (8) as

$$t = h + r + \eta_{h,r,t}, \tag{9}$$

where  $\eta_{h,r,t}$  can be considered an error term for the triple  $\langle h, r, t \rangle$ . Moreover, we define

$$\eta_{\max} = \max_{\langle h,r,t \rangle \in S} \eta_{h,r,t} \tag{10}$$

<sup>6</sup>Created with PyKEEN [1], using 128 epochs, a learning rate of 0.1, the softplus loss function, and default parameters otherwise, as advised by the authors of PyKEEN: <https://github.com/pykeen/pykeen/issues/97>.

<sup>7</sup>This does not mean that TransE does not work. The training data for the very small graph is rather scarce, and two dimensions might not be sufficient to find a good solution here.

Next, we consider two triples  $\langle h_1, r, t \rangle$  and  $\langle h_2, r, t \rangle$ , which share a relation to an object – e.g., in our example, France and Belgium, which both share the relation *partOf* to *EU*. In that case,

$$t = h_1 + r + \eta_{h_1,r,t} \quad (11)$$

and

$$t = h_2 + r + \eta_{h_2,r,t} \quad (12)$$

hold. From that, we get<sup>8</sup>

$$\begin{aligned} h_1 - h_2 &= \eta_{h_2,r,t} - \eta_{h_1,r,t} \\ \Rightarrow |h_1 - h_2| &= |\eta_{h_2,r,t} - \eta_{h_1,r,t}| \\ &= |\eta_{h_2,r,t} + (-\eta_{h_1,r,t})| \\ &\leq |\eta_{h_2,r,t}| + |-\eta_{h_1,r,t}| \\ &= |\eta_{h_2,r,t}| + |\eta_{h_1,r,t}| \\ &\leq 2 \cdot \eta_{\max} \end{aligned} \quad (13)$$

In other words,  $\eta_{\max}$  also imposes an upper bound of two entities sharing a relation to an object. As a consequence, the lower the error in relation prediction, the closer are entities which share a common statement.

This also carries over to entities sharing the same two-hop connection. Consider two further triples  $\langle h_{1a}, r_a, h_1 \rangle$  and  $\langle h_{2a}, r_a, h_2 \rangle$ . In our example, this could be two cities located in the two countries, e.g., *Strasbourg* and *Brussels*. In that case, we would have

$$h_1 = h_{1a} + r_a + \eta_{h_{1a},r_a,h_1} \quad (14)$$

$$h_2 = h_{2a} + r_a + \eta_{h_{2a},r_a,h_2} \quad (15)$$

Substituting this in (11) and (12) yields

$$t = h_{1a} + r_a + \eta_{h_{1a},r_a,h_1} + r + \eta_{h_1,r,t} \quad (16)$$

$$t = h_{2a} + r_a + \eta_{h_{2a},r_a,h_2} + r + \eta_{h_2,r,t}. \quad (17)$$

Consequently, using similar transformations as above, we get

$$\begin{aligned} h_{1a} - h_{2a} &= \eta_{h_{2a},r_a,h_2} - \eta_{h_{1a},r_a,h_1} + \eta_{h_2,r,t} - \eta_{h_1,r,t} \\ \Rightarrow |h_{1a} - h_{2a}| &\leq 4 \cdot \eta_{\max} \end{aligned} \quad (18)$$

Again,  $\eta_{\max}$  constrains the proximity of the two entities  $h_{1a}$  and  $h_{2a}$ , but only half as strictly as for the case of  $h_1$  and  $h_2$ .

#### 4.3. Comparing the two notions of similarity

In the examples above, we can see that embeddings for link prediction have a tendency to project similar instances close to each other in the vector space. Here, the notion of similarity is that two entities are similar if they share a relation to another entity, i.e.,  $e_1$  and  $e_2$  are considered similar if there exist two statements  $\langle e_1, r, t \rangle$  and  $\langle e_2, r, t \rangle$  or

---

<sup>8</sup>Using the triangle inequality for the first inequation.

$\langle h, r, e_1 \rangle$  and  $\langle h, r, e_2 \rangle$ ,<sup>9</sup> or, less strongly, if there exists a chain of such statements. More formally, we can write the notion of similarity between two entities in link prediction approaches as

$$e_1 \approx e_2 \leftarrow \exists t, r : r(e_1, t) \wedge r(e_2, t) \quad (19)$$

$$e_1 \approx e_2 \leftarrow \exists h, r : r(h, e_1) \wedge r(h, e_2) \quad (20)$$

In other words: two entities are similar if they share a common connection to a common third entity.

RDF2vec, on the other hand, covers a wider range of such similarities. Looking at Table 2, we can observe that two entities sharing a common relation to two different objects are also considered similar (*Berlin* and *Mannheim* both share the fact that they are *located in Germany*, hence, their predictions for  $w_{+1}$  and  $w_{+2}$  are similar).

However, there in RDF2vec, similarity can also come in other notions. For example, *Germany* and *USA* are also considered similar, because they both share the relations *headOfGovernment* and *capital*, albeit with different object (i.e., their prediction for  $w_1$  is similar). In contrast, such similarities do not lead to close projections for link prediction embeddings. In fact, in Fig. 9, it can be observed that *USA* and *Germany* are further away than *Germany* and other European countries. In other words, the following two notions of similarity also hold for RDF2vec:

$$e_1 \approx e_2 \leftarrow \exists t_1, t_2, r : r(e_1, t_1) \wedge r(e_2, t_2) \quad (21)$$

$$e_1 \approx e_2 \leftarrow \exists h_1, h_2, r : r(h_1, e_1) \wedge r(h_2, e_2) \quad (22)$$

On a similar argument, RDF2vec also positions entities closer which share *any* relation to another entity. Although this is not visible in the two-dimensional embedding depicted in Fig. 7, RDF2vec would also create vectors with some similarity for *Angela Merkel* and *Berlin*, since they both have a (albeit different) relation to Germany (i.e., their prediction for  $w_{-2}$  is similar). Hence, the following notions of similarity can also be observed in RDF2vec:

$$e_1 \approx e_2 \leftarrow \exists t, r_1, r_2 : r_1(e_1, t) \wedge r_2(e_2, t) \quad (23)$$

$$e_1 \approx e_2 \leftarrow \exists h, r_1, r_2 : r_1(h, e_1) \wedge r_2(h, e_2) \quad (24)$$

The example with *Angela Merkel* and *Berlin* already hints at a slightly different notion of the interpretation of proximity in the vector space evoked by RDF2vec: not only *similar*, but also *related* entities are positioned close in the vector space. This means that to a certain extent, RDF2vec mixes the concepts of similarity and relatedness in its distance function. We will see examples of this in later considerations, and discuss how they interfere with downstream applications.

## 5. Experiments

To compare the two sets of approaches, we use standard setups for evaluating knowledge graph embedding methods for data mining as well as for link prediction.

### 5.1. Experiments on data mining tasks

In our experiments, we follow the setup proposed in [28] and [25]. Those works propose the use of data mining tasks with an external ground truth, e.g., predicting certain indicators or classes for entities. Those entities are then linked to a knowledge graph. Different feature extraction methods – which includes the generation of embedding vectors – can then be compared using a fixed set of learning methods.

The setup of [25] comprises six tasks using 20 datasets in total:

---

<sup>9</sup>The argument in Section 4.2 would also work for shared relations to common heads.

- Five classification tasks, evaluated by accuracy. Those tasks use the same ground truth as the regression tasks (see below), where the numeric prediction target is discretized into high/medium/ low (for the Cities, AAUP, and Forbes dataset) or high/low (for the Albums and Movies datasets). All five tasks are single-label classification tasks.
- Five regression tasks, evaluated by root mean squared error. Those datasets are constructed by acquiring an external target variable for instances in knowledge graphs which is not contained in the knowledge graph per se. Specifically, the ground truth variables for the datasets are: a quality of living indicator for the Cities dataset, obtained from Mercer; average salary of university professors per university, obtained from the AAUP; profitability of companies, obtained from Forbes; average ratings of albums and movies, obtained from Facebook.
- Four clustering tasks (with ground truth clusters), evaluated by accuracy. The clusters are obtained by retrieving entities of different ontology classes from the knowledge graph. The clustering problems range from distinguishing coarser clusters (e.g., cities vs. countries) to finer ones (e.g., basketball teams vs. football teams).
- A document similarity task (where the similarity is assessed by computing the similarity between entities identified in the documents), evaluated by the harmonic mean of Pearson and Spearman correlation coefficients. The dataset is based on the LP50 dataset [18]. It consists of 50 documents, each of which have been annotated with DBpedia entities using DBpedia spotlight [21]. The task is to predict the similarity for each pair of documents.
- An entity relatedness task (where semantic similarity is used as a proxy for semantic relatedness), evaluated by Kendall’s Tau. The dataset is based on the KORE dataset [13]. The dataset consists of 20 seed entities from the YAGO knowledge graph, and 20 related entities each. Those 20 related entities per seed entity have been ranked by humans to capture the strength of relatedness. The task is to rank the entities per seed by relatedness.
- Four semantic analogy tasks (e.g., *Athens is to Greece as Oslo is to X*), which are based on the original datasets on which word2vec was evaluated [22]. The datasets were created by manual annotation. In our evaluation, we aim at predicting the fourth element ( $D$ ) in an analogy  $A : B = C : D$  by considering the closest  $n$  vectors to  $B - A + C$ . If the element is contained the top  $n$  predictions, we consider the answer correct, i.e., the evaluation metric top- $n$  accuracy. In the default setting of the evaluation framework used,  $n$  is set to 2.

Table 3 shows a summary of the characteristics of the datasets used in the evaluation. It can be observed that they cover a wide range of tasks, topics, sizes, and other characteristics (e.g., balance). More details on the construction of the datasets can be found in [25] and [28].

Note that all datasets are provided with predefined instance links to DBpedia. For the smaller ones, the creators of the datasets created and checked the links manually; for the larger ones, the linking had been done heuristically. We used the links provided in the evaluation framework as is, including possible linkage errors.

We follow the evaluation protocol suggested in [25]. This protocol foresees the usage of different algorithms on each task for each embedding (e.g., Naive Bayes, Decision Tree, k-NN, and SVM for classification), and also performs parameter tuning in some cases. In the end, we report the best results per task and embedding method. Those results are depicted in Table 4.

All embeddings are trained on DBpedia 2016-10.<sup>10</sup> For generating the different embedding vectors, we use the DGL-KE framework [47] in the respective standard settings, and we use the RDF2vec vectors provided by the KGvec2go API [26], trained with 500 walks of depth 8 per entity, Skip-Gram, and 200 dimensions. We compare RDF2vec [30], TransE (with L1 and L2 norm) [4], TransR [19], RotatE [37], DistMult [46], RESCAL [23], and ComplEx [40]. To create the embedding vectors with DGL-KE, we use the parameter configurations recommended by the framework, a dimension of 200, and a step maximum of 1,000,000. The RDF2vec<sub>oa</sub> vectors were generated with the same configuration, but using the order-aware variant of Skip-Gram [20,27]. For node2vec, DeepWalk, and KGlove, we use the standard settings and the code provided by the respective authors.<sup>11,12,13</sup> For KGlove, we use the Inverse Predicate Frequency, which has been reported to work well on many tasks by the original paper [7].

<sup>10</sup>The code for the experiments as well as the resulting embeddings can be found at <https://github.com/nheist/KBE-for-Data-Mining>.

<sup>11</sup><https://github.com/D2KLab/entity2vec>

<sup>12</sup><https://github.com/phanein/deepwalk>

<sup>13</sup><https://github.com/miselico/globalRDFEmbeddingsISWC>

Table 3  
Overview on the evaluation datasets

Task	Dataset	# entities	Target variable
Classification	Cities	212	3 classes (67/106/39)
	AAUP	960	3 classes (236/527/197)
	Forbes	1,585	3 classes (738/781/66)
	Albums	1,600	2 classes (800/800)
	Movies	2,000	2 classes (1,000/1,000)
Regression	Cities	212	Numeric [23, 106]
	AAUP	960	Numeric [277, 1009]
	Forbes	1,585	Numeric [0.0, 416.6]
	Albums	1,600	Numeric [15, 97]
	Movies	2,000	Numeric [1, 100]
Clustering	Cities and countries (2k)	4,344	2 clusters (2,000/2,344)
	Cities and countries	11,182	2 clusters (8,838/2,344)
	Cities, countries, albums, movies, AAUP, Forbes	6,357	5 clusters (2,000/960/1,600/212/1,585)
	Teams	4,206	2 clusters (4,185/21)
Document similarity	Pairs of 50 documents with entities	1,225	Numeric similarity score [1.0,5.0]
Entity relatedness	20 × 20 entity pairs	400	Ranking of entities
Semantic analogies	(All) capitals and countries	4,523	Entity prediction
	Capitals and countries	505	Entity prediction
	Cities and states	2,467	Entity prediction
	Countries and currencies	866	Entity prediction

It is noteworthy that the default settings for node2vec and DeepWalk differ in one crucial property. While node2vec interprets the graph as a directed graph by default and only traverses edges in the direction in which they are defined, DeepWalk treats all edges as undirected, i.e., it traverses them in both directions.

From the table, we can observe a few expected and a few unexpected results. First, since RDF2vec is tailored towards classic data mining tasks like classification and regression, it is not much surprising that those tasks are solved better by using RDF2vec (and even slightly better by using RDF2vec<sub>oa</sub>) vectors. Still, some of the link prediction methods (in particular TransE and RESCAL) perform reasonably well on those tasks. In contrast, KGloVe rarely reaches the performance level of RDF2vec, while the two approaches for unlabeled graphs – i.e., DeepWalk and node2vec – behave differently: while the results of DeepWalk are at the lower end of the spectrum, node2vec is competitive. The latter is remarkable, showing that pure neighborhood information, ignoring the direction and edge labels, can be a strong signal when embedding entities.

Referring back to the different notions of similarity that these families of approaches imply (cf. Section 4.3), this behavior can be explained by the tendency of RDF2vec (and also node2vec) to positioning entities closer in the vector space which are more similar to each other (e.g., two cities that are similar). Since it is likely that some of those dimensions are also correlated with the target variable at hand (in other words: they encode some dimension of similarity that can be used to predict the target variable), classifiers and regressors can pick up on those dimensions and exploit them in their prediction model.

What is also remarkable is the performance on the entity relatedness task. While RDF2vec embeddings, as well as node2vec, KGlove, and, to a lesser extent, DeepWalk, reflect entity relatedness to a certain extent, this is not given for any of the link prediction approaches. According to the notions of similarity discussed above, this is reflected in the RDF2vec mechanism: RDF2vec has an incentive to position two entities closer in the vector space if they share relations to a common entity, as shown in Equations (21)-(24). One example is the relatedness of *Apple Inc.* and

Table 4

Results of the different data mining tasks. *DM* denotes approaches originally developed for node representation in data mining, *LP* denotes approaches originally developed for link prediction

Task (metric)	Dataset	RDF2vec (DM)	RDF2vec <sub>OA</sub> (DM)	TransE-L1 (LP)	TransE-L2 (LP)	TransR (LP)	RotatE (LP)	DistMult (LP)	RESCAL (LP)	ComplEx (LP)	node2vec (DM)	DeepWalk (DM)	KGloVe (DM)
Classification (ACC)	AAUP	0.676	0.671	0.628	0.651	0.607	0.617	0.597	0.623	0.602	<b>0.694</b>	0.549	0.558
	Cities	0.810	<b>0.837</b>	0.676	0.752	0.757	0.581	0.666	0.740	0.637	0.774	0.495	0.496
	Forbes	0.610	<b>0.626</b>	0.550	0.601	0.561	0.526	0.601	0.563	0.578	0.618	0.490	0.502
	Albums	0.774	0.787	0.637	0.746	0.728	0.550	0.666	0.678	0.693	<b>0.789</b>	0.543	0.548
	Movies	0.739	0.736	0.603	0.728	0.715	0.567	0.668	0.693	0.655	<b>0.763</b>	0.555	0.563
Clustering (ACC)	Cities and countries (2K)	0.758	0.931	0.982	<b>0.994</b>	0.962	0.510	0.957	0.991	0.955	0.939	0.557	0.623
	Cities and countries	0.696	0.760	0.953	0.979	0.952	0.691	0.909	<b>0.990</b>	0.591	0.743	0.817	0.765
	Cities, albums, movies, AAUP, Forbes	0.926	0.928	<b>0.946</b>	0.944	0.908	0.860	0.878	0.936	0.914	0.930	0.335	0.525
	Teams	0.917	0.958	0.887	<b>0.977</b>	0.844	0.853	0.883	0.881	0.881	0.931	0.830	0.740
Regression (RMSE)	AAUP	68.745	<b>66.505</b>	81.503	69.728	88.751	80.177	78.337	72.880	73.665	68.007	103.235	98.794
	Cities	15.601	<b>13.486</b>	19.694	14.455	13.558	26.846	19.785	15.137	19.809	15.363	25.323	24.151
	Forbes	36.459	36.124	37.589	38.398	39.803	38.343	38.037	<b>35.489</b>	37.877	35.684	41.384	40.141
	Albums	11.930	<b>11.597</b>	14.128	12.589	12.789	14.890	13.452	13.537	13.009	15.165	15.129	11.739
	Movies	19.648	<b>11.739</b>	23.286	20.635	20.699	23.878	22.161	21.362	22.229	18.877	24.215	24.000
Semantic analogies (precision@k)	(All) capitals and countries	0.685	0.895	0.709	0.675	<b>0.938</b>	0.377	0.782	0.211	0.814	0.284	0.000	0.011
	Capitals and countries	0.648	0.913	0.840	0.792	<b>0.937</b>	0.640	0.802	0.312	0.864	0.164	0.000	0.043
	Cities and state	0.342	<b>0.629</b>	0.335	0.209	0.392	0.294	0.379	0.089	0.309	0.068	0.000	0.029
	Currency (and countries)	0.339	<b>0.427</b>	0.005	0.285	0.143	0.000	0.001	0.000	0.000	0.420	0.005	0.003
Document similarity (harmonic mean)	LP50	0.348	0.307	0.343	0.397	<b>0.434</b>	0.326	0.360	0.344	0.341	0.333	0.243	0.225
Entity relatedness (KendallTau)	KORE	0.504	<b>0.779</b>	0.002	-0.081	0.139	-0.039	0.147	0.087	0.115	0.525	0.129	0.421



*Steve Jobs* – here, we can observe the two statements

*product(Apple Inc., iPhone)*

*knownfor(Steve Jobs, iPhone)*

in DBpedia, among others. Those lead to similar vectors in RDF2vec according to Equation (23). A similar argument can be made for node2vec and DeepWalk, and also for KGlove, which looks at global cooccurrences of entities, i.e., it also favors closer embeddings of related entities.

The same behavior of RDF2vec – i.e., assigning close vectors to *related* entities – also explains the comparatively bad results of RDF2vec on the first two clustering tasks. Here, the task is to separate cities and countries in two clusters, but since a city is also related to the country it is located in, RDF2vec may position that city and country rather closely together (RDF2vec<sub>oa</sub> changes that behavior, as argued in [27], and hence produces better results for the clustering problems). Hence, that city has a certain probability of ending up in the same cluster as the country. The latter two clustering tasks are different: the third one contains five clusters (cities, albums, movies, universities, and companies), which are less likely to be strongly related (except universities and companies to cities) and therefore are more likely to be projected in different areas in the vector space. Here, the difference of RDF2vec to the best performing approaches (i.e., TransE-L1 and TransE-L2) is not that severe. The same behavior can also be observed for the other embedding approaches for data mining, i.e., node2vec, DeepWalk, and KGlove, which behave similarly in that respect.

The problem of relatedness being mixed with similarity does not occur so strongly for *homogeneous* sets of entities, as in the classification and regression tasks, where all entities are of the same kind (cities, companies, etc.) – here, two companies which are related (e.g., because one is a holding of the other) can also be considered similar to a certain degree (in that case, they are both operating in the same branch). This also explains why the fourth clustering task (where the task is to assign sports teams to clusters by the type of sports) works well for RDF2vec – here, the entities are again homogeneous.

At the same time, the test case of clustering teams can also be used to explain why link prediction approaches work well for that kind of tasks: here, it is likely that two teams in the same sports share a relation to a common entity, i.e., they fulfill Equations (19) and (20). Examples include participation in the same tournaments or common former players.

The semantic analogies task also reveals some interesting findings. First, it should be noted that the relations which form the respective analogies (capital, state, and currency) is contained in the knowledge graph used for the computation. That being said, we can see that most of the link prediction results (except for RotatE and RESCAL) perform reasonably well here. Particularly, the first cases (capitals and countries) can be solved particularly well in those cases, as this is a 1:1 relation, which is the case in which link prediction is a fairly simple task. On the other hand, most of the data-mining-centric approaches (i.e., node2vec, DeepWalk, KGlove) solve this problem relatively bad. A possible explanation is that the respective entities belong to the strongly interconnected head entities of the knowledge graphs, and also the false solutions are fairly close to each other in the graph (e.g., US Dollar and Euro are interconnected through various short paths). This makes it hard for approaches concentrating on a common neighborhood to produce decent results here.

On the other hand, the currency case is solved particularly bad by most of the link prediction results. This relation is an n:m relation (there are countries with more than one official, unofficial, or historic currency, and many currencies, like the Euro, are used across many countries. Moreover, looking into DBpedia, this relation contains a lot of mixed usage and is not maintained with very high quality. For example, DBpedia lists 33 entities whose currency is US Dollars<sup>14</sup> – the list contains historic entities (e.g., West Berlin), errors (e.g., Netherlands), and entities which are not countries (e.g., OPEC), but the United States are not among those. For such kind of relations which contain a certain amount of noise and heterogeneous information, many link prediction approaches are obviously not well suited.

---

<sup>14</sup>[http://dbpedia.org/page/United\\_States\\_dollar](http://dbpedia.org/page/United_States_dollar)

RDF2vec, in contrast, can deal reasonably well with that case. Here, two effects interplay when solving such tasks: (i) as shown above, relations are encoded by the proximity in RDF2vec to a certain extent, i.e., the properties in Equations (3) and (4) allow to perform analogy reasoning in the RDF2vec space in general. Moreover, (ii) we have already seen the tendency of RDF2vec to position *related* entities in relative proximity. Thus, for RDF2vec, it can be assumed that the following holds:

$$UK \approx \text{Pound Sterling} \quad (25)$$

$$USA \approx \text{US Dollar} \quad (26)$$

Since we can rephrase the first equation as

$$\text{Pound Sterling} - UK \approx 0 \quad (27)$$

we can conclude that analogy reasoning in RDF2vec would yield

$$\text{Pound Sterling} - UK + USA \approx \text{US Dollar} \quad (28)$$

Hence, in RDF2vec, two effects – the preservation of relation vectors as well as the proximity of related entities – are helpful for analogy reasoning, and the two effects also work for rather noisy cases. However, for cases which are 1:1 relations in the knowledge graph with rather clean training data available, link prediction approaches are better suited for analogy reasoning.

## 5.2. Experiments on link prediction tasks

In a second series of experiments, we analyze if we can use embedding methods developed for similarity computation, like RDF2vec, also for link prediction. We use the two established tasks WN18 and FB15k for a comparative study.

While link prediction methods are developed for the task at hand, approaches developed for data mining are not. Although RDF2vec computes vectors for relations, they do not necessarily follow the same notion as relation vectors for link prediction, as discussed above. Hence, we investigate two approaches:

1. We average the difference for each pair of a head and a tail for each relation  $r$ , and use that as average as a proxy for a relation vector for prediction, as shown in Equation (4). The predictions are the entities whose embedding vectors are the closest to the approximate prediction. This method is denoted as *avg*.
2. For predicting the tail of a relation, we train a neural network to predict an embedding vector of the tail based embedding vectors, as shown in Fig. 10. The predictions for a triple  $\langle h, r, ? \rangle$  are the entities whose embedding vectors are closest to the predicted vector for  $h$  and  $r$ . As similar network is trained to predict  $h$  from  $r$  and  $t$ . This method is denoted as *ANN*.

We trained the RDF2vec embeddings with 2,000 walks, a depth of 4, a dimension of 200, a window of 5, and 25 epochs in SG mode. For the second prediction approach, the two neural networks each use two hidden layers of size 200, and we use 15 epochs, a batch size of 1,000, and mean squared error as loss. KGlove, node2vec, and DeepWalk do not produce any vectors for relations. Hence, we only use the *avg* strategy for those approaches.

The results of the link prediction experiments are shown in Table 5.<sup>15</sup> We can observe that the RDF2vec based approaches perform at the lower end of the spectrum. The *avg* approach outperforms DistMult and RESCAL on WN18, and both approaches are about en par with RESCAL on FB15k. Except for node2vec on FB15k, the other data mining approaches fail at producing sensible results.

While the results are not overwhelming, they show that similarity of entities, as RDF2vec models it, is at least a useful signal for implementing a link prediction approach.

<sup>15</sup>The code for the experiments can be found at [https://github.com/janothan/kbc\\_rdf2vec](https://github.com/janothan/kbc_rdf2vec).

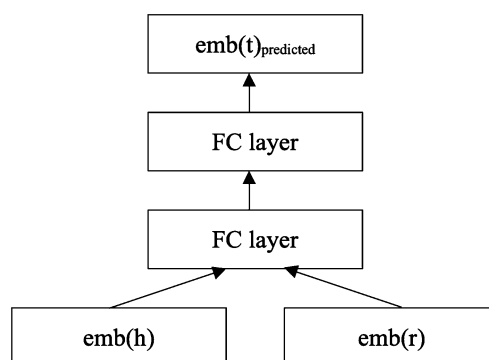


Fig. 10. Training a neural network for link prediction with RDF2vec.

### 5.3. Discussion

As already discussed above, the notion of similarity which is conveyed by RDF2vec mixes *similarity* and *relatedness*. This can be observed, e.g., when querying for the 10 closest concepts to *Angela Merkel* (the chancellor, i.e., head of government in Germany) in DBpedia in the different spaces, as shown in Table 6. The approach shows a few interesting effects:

- While most of the approaches (except for RotatE, KGlove and DeepWalk) provide a clean list of people, RDF2vec brings up a larger variety of results, containing also *Germany* and *Berlin* (and also a few results which are not instances, but relations; however, those could be filtered out easily in downstream applications if necessary). This demonstrates the property of RDF2vec of mixing *similarity* and *relatedness*. The people in the RDF2vec result set are all related to Angela Merkel: Joachim Gauck was president during her chancellorship, Norbert Lammert was the head of parliament, Stanislaw Tillich was a leading board member in the same party as Merkel, and Andreas Voßkuhle was the head of the highest court during her chancellorship.
- The approaches at hand have different foci in determining similarity. For example, TransE-L1 outputs mostly German politicians (Schröder, Gauck, Trittin, Gabriel, Westerwelle, Wulff) and former presidents of other countries (Buchanan as a former US president, Sarkozy and Chirac as former French presidents) TransE-L2 outputs a list containing many former German chancellors (Schröder, Kohl, Adenauer, Schmidt, Kiesinger, Erhardt), TransR mostly lists German party leaders (Gabriel, Steinmeier, Rösler Schröder, Wulff, Westerwelle, Kohl, Trittin). Likewise, node2vec produces a list of German politicians, with the exception of Merkel’s husband Joachim Sauer.<sup>16</sup> In all of those cases, the persons share some property with the query entity *Angela Merkel* (profession, role, nationality, etc.), but similarity is usually affected only by one of those properties. In other words: one notion of similarity *dominates* the others.
- In contrast, the persons in the output list of RDF2vec are *related to the query entity in different* respects. In particular, they played different roles during Angela Merkel’s chancellorship (Gauck was the German president, Lammert was the chairman of the parliament, and Voßkuhle was the chairman of the federal court). Here, there is no *dominant* property, instead, similarity (or rather: relatedness) is encoded along various properties. RDF2vec<sub>oa</sub> brings up a results which is slightly closer to the politicians lists of the other approaches, while the result list of KGlove looks more like a random list of entities. A similar observation can be made for DeepWalk, which, with the exception of the first result (which is a German politician) does not produce any results seemingly related to the query concept at hand.

With that observation in mind, we can come up with an initial set of recommendations for choosing embedding approaches:

<sup>16</sup>The remaining approaches – RotatE, DistMult, RESCAL, ComplEx, KGlove, DeepWalk – produce lists of (mostly) persons which, in their majority, share no close link to the query concept Angela Merkel.

Table 5

Results of the link prediction tasks on WN18 and FB15K. Results for TransE and RESCAL from [4], results for RotatE from [37], results for DistMult from [46], results for TransR from [19]. *DM* denotes approaches originally developed for node representation in data mining, *LP* denotes approaches originally developed for link prediction

Dataset	Metric	RDF2vec (AVG) (DM)	RDF2vec (ANN) (DM)	RDF2vec <sub>OA</sub> (AVG) (DM)	RDF2vec <sub>OA</sub> (ANN) (DM)	TransE (LP)	TransR (LP)	RotatE (LP)	DistMult (LP)	RESCAL (LP)	ComplEx (LP)	node2vec (AVG) DM	DeepWalk (AVG) DM	KGlove (AVG) (DM)
WN18	Mean rank raw	147	353	64	77	263	232	-	-	1180	-	17	6112	8247
	Mean rank filtered	135	342	53	66	251	219	309	-	1163	-	10	6106	8243
	HITS@10 raw	64.4	49.7	65.9	66.6	75.4	78.3	-	-	37.2	-	12.3	5.7	1.7
	HITS@10 filtered	71.3	55.4	73.0	75.08	89.2	91.7	<b>95.9</b>	57.7	52.8	94.7	14.2	6.0	1.7
FB15K	Mean rank raw	399	349	168	443	243	226	-	-	828	-	192	2985	2123
	Mean rank filtered	347	303	120	90	125	78	40	-	683	-	138	2939	2077
	HITS@10 raw	35.3	34.3	47.9	30.9	34.9	43.8	-	-	28.4	-	44.7	7.2	11.1
	HITS@10 filtered	40.5	41.8	56.5	37.4	47.1	65.5	88.4	<b>94.2</b>	44.1	84.0	53.3	7.7	11.1

Table 6  
Closest concepts to *Angela Merkel* in the different embedding approaches used

RDF2vec	TransE-L1	TransE-L2	TransR
Joachim Gauck	Gerhard Schröder	Gerhard Schröder	Sigmar Gabriel
Norbert Lammert	James Buchanan	Helmut Kohl	Frank-Walter Steinmeier
Stanislaw Tillich	Neil Kinnock	Konrad Adenauer	Philipp Rösler
Andreas Voßkuhle	Nicolas Sarkozy	Helmut Schmidt	Gerhard Schröder
Berlin	Joachim Gauck	Werner Faymann	Joachim Gauck
German language	Jacques Chirac	Alfred Gusenbauer	Christian Wulff
Germany	Jürgen Trittin	Kurt Georg Kiesinger	Guido Westerwelle
federalState	Sigmar Gabriel	Philipp Scheidemann	Helmut Kohl
Social Democratic Party	Guido Westerwelle	Ludwig Erhard	Jürgen Trittin
deputy	Christian Wulff	Wilhelm Marx	Jens Böhrens
RotatE	DistMult	RESCAL	Complex
Pontine raphe nucleus	Gerhard Schröder	Gerhard Schröder	Gerhard Schröder
Jonathan W. Bailey	Milan Truban	Kurt Georg Kiesinger	Diána Mészáros
Zokwang Trading	Maud Cuney Hare	Helmut Kohl	Francis M. Bator
Steven Hill	Tristan Matthiae	Annemarie Huber-Hotz	William B. Bridges
Chad Kreuter	Gerda Hasselfeldt	Wang Zhaoguo	Mette Vestergaard
Fred Hibbard	Faustino Sainz Muñoz	Franz Vranitzky	Ivan Rosenqvist
Mallory Ervin	Joachim Gauck	Bogdan Klich	Edward Clouston
Paulinho Kobayashi	Carsten Linnemann	İrsen Küçük	Antonio Capuzzi
Fullmetal Alchemist and the Broken Angel	Norbert Blüm	Helmut Schmidt	Steven J. McAuliffe
Archbishop Dorotheus of Athens	Neil Hood	Mao Zedong	Jenkin Coles
KGloVe	RDF2vec OA	node2vec	DeepWalk
Aurora Memorial National Park	Joachim Gauck	Sigmar Gabriel	Manuela Schwesig
Lithuanian Wikipedia	Norbert Lammert	Guido Westerwelle	Irwin Fridovich
Baltic states	Stanislaw Tillich	Christian Wulff	Holstein Kiel Dominik Schmidt
The Monarch (production team)	Andreas Voßkuhle	Jürgen Trittin	Ella Germein
Leeds Ladies F.C. Lauryn Colman	Berlin	Wolfgang Schäuble	Goyang Citizen FC Do Sang-Jin
Steven Marković	German language	Joachim Gauck	Sean Cashman
Funk This (George Porter Jr. album)	Germany	Philipp Rösler	Chia Chiao
A Perfect Match (Ella Fitzgerald album)	Christian Wulff	Joachim Sauer	Albrix Niigata Goson Sakai
Salty liquorice	Gerhard Schröder	Franz Müntefering	Roz Kelly
WMMU-FM	federalState	Frank-Walter Steinmeier	Alberto Penny

- Approaches for data mining (RDF2vec, KGlove, node2vec, and DeepWalk) work well when dealing with sets of *homogeneous entities*. Here, the problem of confusing related entities (like Merkel and Berlin) is negligible, because all entities are of the same kind anyways. In those cases, RDF2vec captures the finer distinctions between the entities better than embeddings for link prediction, and it encodes a larger variety of semantic relations.
- From the approaches for data mining, those which respect order (RDF2vec<sub>oa</sub> and node2vec) work better than those which do not (classic RDF2vec, KGlove, and DeepWalk).<sup>17</sup>
- For problems where *heterogeneous sets of entities* are involved, embeddings for link prediction often do a better job in telling different entities apart.

Link prediction is a problem of the latter kind: in embedding spaces where different types are properly separated, link prediction mistakes are much rarer. Given an embedding space where entities of the same type are always closer

<sup>17</sup>As discussed above, this comments holds for the *default* configuration of node2vec and DeepWalk used in this paper.

than entities of a different type, a link prediction approach will always rank all “compatible” entities higher than all incompatible ones. Consider the following example in FB15k:

*instrument(Gil Scott Heron, ?)*

Here, music instruments are expected in the object position. However, approaches tailored towards capturing node similarity, e.g., classic RDF2vec, will suggest among plausible candidates such as *electric guitar* and *acoustic guitar*, also *guitarist* and *Jimmy Page* (who is a well-known guitarist). While *electric guitar*, *guitarist*, and *Jimmy Page* are semantically related, not all of them are sensible predictions here, and the fact that RDF2vec reflects that semantic relatedness is a drawback in link prediction.

The same argument underlies an observation made by Zouaq and Martel [48]: the authors found that RDF2vec is particularly well suited for distinguishing fine-grained entity classes (as opposed to coarse-grained entity classification). For fine-grained classification (e.g., distinguishing guitar players from singers), all entities to be classified are already of the same coarse class (e.g., musician), and RDF2vec is very well suited for capturing the finer differences. However, for coarse classifications, misclassifications by mistaking relatedness for similarity become more salient.

From the observations made in the link prediction task, we can come up with another recommendation:

- For relations which come with rather clean data quality, link prediction approaches work well. However, for more noisy data, RDF2vec has a higher tendency of creating useful embedding vectors.

For the moment, this is a hypothesis, which should be hardened, e.g., by performing controlled experiments on artificially noised link prediction tasks.

## 6. Conclusion and outlook

In this paper, we have compared two use cases and families of knowledge graph embeddings which have, up to today, not undergone any thorough direct comparison: approaches developed for data mining, such as RDF2vec, and approaches developed for link prediction, such as TransE and its descendants.

We have argued that the two approaches actually do something similar, albeit being designed with different goals in mind. To support this argument, we have run two sets of experiments which examined how well the different approaches work if applied in the respective other setup. We show that, to a certain extent, embedding approaches designed for link prediction can be applied in data mining and vice versa, however, there are differences in the outcome.

From the experiments, we have also seen that proximity in the embedding spaces works differently for the two families of approaches: in RDF2vec, proximity encodes both similarity and relatedness, while TransE and its descendants rather encode similarity alone. On the other hand, for entities that are of the same type, RDF2vec covers finer-grained similarities better. Moreover, RDF2vec seems to work more stably in cases where the knowledge graphs are rather noisy and weakly adherent to their schema.

These findings give rise both for a recommendation and some future work. First, in use cases where relatedness plays a role next to similarity, or in use cases where all entities are of the same type, approaches like RDF2vec may yield better results. On the other hand, for cases with mixed entity types where it is important to separate the types, link prediction embeddings might yield better results.

Since the set of knowledge graphs used in our experiments is limited, we can, however, not come up with recommendations of which kind of embedding is better suited for which kind of knowledge graph. While we expect that there are differences with respect to different characteristics of the graph – e.g., homogeneity, link degree and cardinality distributions, density and sparsity, schema size and variety – both theoretical considerations and experimental evaluations in that direction are subject to future work.

Moreover, the open question remains whether it is possible to develop embedding methods that combine the best of both worlds – e.g., that provide both the coarse type separation of TransE and its descendants and the fine type separation of RDF2vec, or that support competitive link prediction while also representing relatedness. We expect to see some interesting developments along these lines in the future.

## Acknowledgement

The publication of this article was funded by the Ministry of Science, Research and the Arts Baden-Württemberg and the University of Mannheim.

## References

- [1] M. Ali, M. Berrendorf, C.T. Hoyt, L. Vermue, S. Sharifzadeh, V. Tresp and J. Lehmann, PyKEEN 1.0: A Python library for training and evaluating knowledge graph embeddings, *Journal of Machine Learning Research* **22**(82) (2021), 1–6.
- [2] F. Bakhshandegan Moghaddam, C. Draschner, J. Lehmann and H. Jabeen, Literal2Feature: An automatic scalable RDF graph feature extractor, in: *Further with Knowledge Graphs*, IOS Press, 2021, pp. 74–88. doi:10.3233/SSW210036.
- [3] S. Basu, S. Chakraborty, A. Hassan, S. Siddique and A. Anand, ERLKG: Entity representation learning and knowledge graph based association analysis of COVID-19 through mining of unstructured biomedical corpora, in: *Proceedings of the First Workshop on Scholarly Document Processing*, 2020, pp. 127–137. doi:10.18653/v1/2020.sdp-1.15.
- [4] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston and O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Advances in Neural Information Processing Systems*, 2013, pp. 2787–2795.
- [5] R. Celebi, H. Uyar, E. Yasar, O. Gumus, O. Dikenelli and M. Dumontier, Evaluation of knowledge graph embedding approaches for drug-drug interaction prediction in realistic settings, *BMC Bioinformatics* **20**(1) (2019), 1–14. doi:10.1186/s12859-019-3284-5.
- [6] J. Chen, X. Chen, I. Horrocks, E.B. Myklebust and E. Jimenez-Ruiz, Correcting knowledge base assertions, in: *Proceedings of the Web Conference 2020*, 2020, pp. 1537–1547. doi:10.1145/3366423.3380226.
- [7] M. Cochez, P. Ristoski, S.P. Ponzetto and H. Paulheim, Global RDF vector space embeddings, in: *International Semantic Web Conference*, Springer, 2017, pp. 190–207. doi:10.1007/978-3-319-68288-4\_12.
- [8] Y. Dai, S. Wang, N.N. Xiong and W. Guo, A survey on knowledge graph embedding: Approaches, applications and benchmarks, *Electronics* **9**(5) (2020), 750. doi:10.3390/electronics9050750.
- [9] D. Daza, M. Cochez and P. Groth, Inductive entity representations from text via link prediction, in: *Proceedings of the Web Conference 2021*, 2021, pp. 798–808. doi:10.1145/3442381.3450141.
- [10] T. Dettmers, P. Minervini, P. Stenetorp and S. Riedel, Convolutional 2d knowledge graph embeddings, in: *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [11] 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* **12**(4) (2021), 617–647. doi:10.3233/SW-200404.
- [12] X. Han, S. Cao, X. Lv, Y. Lin, Z. Liu, M. Sun and J. Li, OpenKE: An open toolkit for knowledge embedding, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 139–144. doi:10.18653/v1/D18-2024.
- [13] J. Hoffart, S. Seufert, D.B. Nguyen, M. Theobald and G. Weikum, KORE: Keyphrase overlap relatedness for entity disambiguation, in: *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, 2012, pp. 545–554. doi:10.1145/2396761.2396832.
- [14] S. Ji, S. Pan, E. Cambria, P. Marttinen and S.Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE Transactions on Neural Networks and Learning Systems* (2021). doi:10.1109/TNNLS.2021.3070843.
- [15] M.R. Karim, M. Cochez, J.B. Jares, M. Uddin, O. Beyan and S. Decker, Drug-drug interaction prediction based on knowledge graph embeddings and convolutional-LSTM network, in: *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, 2019, pp. 113–123. doi:10.1145/3307339.3342161.
- [16] M. Kejriwal and P. Szekely, Supervised typing of big graphs using semantic embeddings, in: *Proceedings of the International Workshop on Semantic Big Data*, 2017, pp. 1–6. doi:10.1145/3066911.3066918.
- [17] N. Lavrač, B. Škrlj and M. Robnik-Šikonja, Propositionalization and embeddings: Two sides of the same coin, *Machine Learning* **109**(7) (2020), 1465–1507. doi:10.1007/s10994-020-05890-8.
- [18] M.D. Lee, B. Pincombe and M. Welsh, An empirical evaluation of models of text document similarity, in: *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 27, 2005, <https://hdl.handle.net/2440/28910>.
- [19] Y. Lin, Z. Liu, M. Sun, Y. Liu and X. Zhu, Learning entity and relation embeddings for knowledge graph completion, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29, 2015, pp. 2181–2187.
- [20] W. Ling, C. Dyer, A.W. Black and I. Trancoso, Two/too simple adaptations of word2vec for syntax problems, in: *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2015, pp. 1299–1304. doi:10.3115/v1/N15-1142.
- [21] P.N. Mendes, M. Jakob, A. García-Silva and C. Bizer, DBpedia spotlight: Shedding light on the web of documents, in: *Proceedings of the 7th International Conference on Semantic Systems*, 2011, pp. 1–8. doi:10.1145/2063518.2063519.
- [22] T. Mikolov, K. Chen, G. Corrado and J. Dean, Efficient estimation of word representations in vector space, in: *International Conference on Learning Representations*, 2013.
- [23] M. Nickel, V. Tresp and H.-P. Kriegel, A three-way model for collective learning on multi-relational data, in: *International Conference on Machine Learning*, 2011, pp. 809–816.

- [24] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic Web* **8**(3) (2017), 489–508. doi:[10.3233/SW-160218](https://doi.org/10.3233/SW-160218).
- [25] M.A. Pellegrino, A. Altalba, M. Garofalo, P. Ristoski and M. Cochez, GEval: A modular and extensible evaluation framework for graph embedding techniques, in: *European Semantic Web Conference*, Springer, 2020, pp. 565–582. doi:[10.1007/978-3-030-49461-2\\_33](https://doi.org/10.1007/978-3-030-49461-2_33).
- [26] J. Portisch, M. Hladik and H. Paulheim, KGvec2go – knowledge graph embeddings as a service, in: *Proceedings of the 12th Language Resources and Evaluation Conference*, 2020, pp. 5641–5647, <https://aclanthology.org/2020.lrec-1.692>.
- [27] J. Portisch and H. Paulheim, Putting RDF2vec in order, in: *International Semantic Web Conference, Posters and Demos*, 2021.
- [28] P. Ristoski, G.K.D. De Vries and H. Paulheim, A collection of benchmark datasets for systematic evaluations of machine learning on the semantic web, in: *International Semantic Web Conference*, Springer, 2016, pp. 186–194. doi:[10.1007/978-3-319-46547-0\\_20](https://doi.org/10.1007/978-3-319-46547-0_20).
- [29] P. Ristoski and H. Paulheim, A comparison of propositionalization strategies for creating features from linked open data, *Linked Data for Knowledge Discovery* **6** (2014).
- [30] P. Ristoski and H. Paulheim, RDF2vec: RDF graph embeddings for data mining, in: *International Semantic Web Conference*, Springer, 2016, pp. 498–514. doi:[10.1007/978-3-319-46523-4\\_30](https://doi.org/10.1007/978-3-319-46523-4_30).
- [31] P. Ristoski and H. Paulheim, Semantic web in data mining and knowledge discovery: A comprehensive survey, *Journal of Web Semantics* **36** (2016), 1–22. doi:[10.1016/j.websem.2016.01.001](https://doi.org/10.1016/j.websem.2016.01.001).
- [32] P. Ristoski, J. Rosati, T. Di Noia, R. De Leone and H. Paulheim, RDF2Vec: RDF graph embeddings and their applications, *Semantic Web* **10**(4) (2019), 721–752. doi:[10.3233/SW-180317](https://doi.org/10.3233/SW-180317).
- [33] A. Rossi, D. Barbosa, D. Firmani, A. Marinata and P. Meriardo, Knowledge graph embedding for link prediction: A comparative analysis, *ACM Transactions on Knowledge Discovery from Data (TKDD)* **15**(2) (2021), 1–49. doi:[10.1145/3424672](https://doi.org/10.1145/3424672).
- [34] B. Shi and T. Weninger, Open-world knowledge graph completion, in: *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [35] R. Sofronova, R. Biswas, M. Alam and H. Sack, Entity typing based on RDF2vec using supervised and unsupervised methods, in: *European Semantic Web Conference*, Springer, 2020, pp. 203–207. doi:[10.1007/978-3-030-62327-2\\_35](https://doi.org/10.1007/978-3-030-62327-2_35).
- [36] B. Steenwinckel, G. Vandewiele, I. Rausch, P. Heyvaert, R. Taelman, P. Colpaert, P. Simoens, A. Dimou, F. De Turck and F. Ongenaë, Facilitating the analysis of COVID-19 literature through a knowledge graph, in: *International Semantic Web Conference*, Springer, 2020, pp. 344–357. doi:[10.1007/978-3-030-62466-8\\_22](https://doi.org/10.1007/978-3-030-62466-8_22).
- [37] Z. Sun, Z.-H. Deng, J.-Y. Nie and J. Tang, RotatE: Knowledge graph embedding by relational rotation in complex space, in: *International Conference on Learning Representations*, 2018.
- [38] P.-N. Tan, M. Steinbach and V. Kumar, *Introduction to Data Mining*, Pearson Education India, 2016.
- [39] K. Toutanova, D. Chen, P. Pantel, H. Poon, P. Choudhury and M. Gamon, Representing text for joint embedding of text and knowledge bases, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 1499–1509. doi:[10.18653/v1/D15-1174](https://doi.org/10.18653/v1/D15-1174).
- [40] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier and G. Bouchard, Complex embeddings for simple link prediction, in: *International Conference on Machine Learning*, PMLR, 2016, pp. 2071–2080.
- [41] G. Vandewiele, B. Steenwinckel, M. Weyns, P. Bonte, F. Ongenaë and F. De Turck, pyRDF2Vec: A python library for RDF2Vec, 2020, <https://github.com/IBCNServices/pyRDF2Vec>.
- [42] Q. Wang, Z. Mao, B. Wang and L. Guo, Knowledge graph embedding: A survey of approaches and applications, *IEEE Transactions on Knowledge and Data Engineering* **29**(12) (2017), 2724–2743. doi:[10.1109/TKDE.2017.2754499](https://doi.org/10.1109/TKDE.2017.2754499).
- [43] X. Wang, T. Gao, Z. Zhu, Z. Zhang, Z. Liu, J. Li and J. Tang, KEPLER: A unified model for knowledge embedding and pre-trained language representation, *Transactions of the Association for Computational Linguistics* **9** (2021), 176–194. doi:[10.1162/tacl\\_a\\_00360](https://doi.org/10.1162/tacl_a_00360).
- [44] Y. Wang, L. Dong, X. Jiang, X. Ma, Y. Li and H. Zhang, KG2Vec: A node2vec-based vectorization model for knowledge graph, *Plos One* **16**(3) (2021), e0248552. doi:[10.1371/journal.pone.0248552](https://doi.org/10.1371/journal.pone.0248552).
- [45] R. Xie, Z. Liu, J. Jia, H. Luan and M. Sun, Representation learning of knowledge graphs with entity descriptions, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30, 2016.
- [46] B. Yang, W. Yih, X. He, J. Gao and L. Deng, Embedding entities and relations for learning and inference in knowledge bases, in: *International Conference on Learning Representations*, 2015.
- [47] D. Zheng, X. Song, C. Ma, Z. Tan, Z. Ye, J. Dong, H. Xiong, Z. Zhang and G. Karypis, DGL-KE: Training knowledge graph embeddings at scale, in: *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 739–748. doi:[10.1145/3397271.3401172](https://doi.org/10.1145/3397271.3401172).
- [48] A. Zouaq and F. Martel, What is the schema of your knowledge graph? Leveraging knowledge graph embeddings and clustering for expressive taxonomy learning, in: *Proceedings of the International Workshop on Semantic Big Data*, 2020, pp. 1–6. doi:[10.1145/3391274.3393637](https://doi.org/10.1145/3391274.3393637).