

RDF2Vec: RDF Graph Embeddings and Their Applications

Petar Ristoski^a, Jessica Rosati^{b,c}, Tommaso Di Noia^b, Renato De Leone^c, Heiko Paulheim^a

^a *Data and Web Science Group, University of Mannheim, B6, 26, 68159 Mannheim*

E-mail: petar.ristoski@informatik.uni-mannheim.de, heiko@informatik.uni-mannheim.de

^b *Polytechnic University of Bari, Via Orabona, 4, 70125 Bari*

E-mail: jessica.rosati@poliba.it, tommaso.dinoia@poliba.it

^c *University of Camerino, Piazza Cavour 19/f, 62032 Camerino*

E-mail: jessica.rosati@unicam.it, renato.deleone@unicam.it

Abstract. Linked Open Data has been recognized as a valuable source for background information in many data mining and information retrieval tasks. However, most of the existing tools require features in propositional form, i.e., a vector of nominal or numerical features associated with an instance, while Linked Open Data sources are graphs by nature. In this paper, we present RDF2Vec, an approach that uses language modeling approaches for unsupervised feature extraction from sequences of words, and adapts them to RDF graphs. We generate sequences by leveraging local information from graph sub-structures, harvested by Weisfeiler-Lehman Subtree RDF Graph Kernels and graph walks, and learn latent numerical representations of entities in RDF graphs. We evaluate our approach on three different tasks: (i) standard machine-learning tasks (ii) entity and document modeling (iii) content-based recommender systems. The evaluation shows that the proposed entity embeddings outperform existing techniques, and that feature vector representations of general knowledge graphs such as DBpedia and Wikidata can be easily reused for different tasks.

Keywords: Graph Embeddings, Linked Open Data, Data Mining, Document Semantic Similarity, Entity Relatedness, Recommender Systems

1. Introduction

Since its launching, the Linked Open Data (LOD) [67] initiative has played a leading role in the rise of a new breed of open and interlinked knowledge bases freely accessible on the Web, each of them being part of a huge decentralized data space, the LOD cloud. This latter is implemented as an open, inter-linked collection of datasets in machine-interpretable form, mainly built on top of two *World Wide Web Consortium* (W3C) standards: RDF¹ and SPARQL². Currently, the LOD cloud consists of about 1,000 inter-linked datasets covering multiple domains from life

science to government data [67]. The LOD cloud has been recognized as a valuable source of background knowledge in data mining and knowledge discovery in general [62], as well as for information retrieval and recommender systems [10]. Augmenting a dataset with features taken from Linked Open Data can, in many cases, improve the results of the problem at hand, while externalizing the cost of maintaining that background knowledge [49].

Most algorithms work with a propositional *feature vector* representation of the data, which means that each instance is represented as a vector of features $\langle f_1, f_2, \dots, f_n \rangle$, where the features are either binary (i.e., $f_i \in \{true, false\}$), numerical (i.e., $f_i \in \mathbb{R}$), or nominal (i.e., $f_i \in S$, where S is a finite set of symbols). Linked Open Data, however, comes in the form of *graphs*, connecting resources with types and

¹<http://www.w3.org/TR/2004/>

REC-rdf-concepts-20040210/, 2004.

²<http://www.w3.org/TR/rdf-sparql-query/>, 2008

relations, backed by a schema or ontology. In order to make LOD accessible to existing data mining tools, an initial *propositionalization* [28] of the corresponding graph is required. Even though the new set of propositional features do not encode all the knowledge available in the original ontological data, they can be effectively used to train a model via machine learning techniques and algorithms. Usually, binary features (e.g., `true` if a type or relation exists, `false` otherwise) or numerical features (e.g., counting the number of relations of a certain type) are used [52,60]. Other variants, e.g., counting different graph sub-structures have also been proposed and used [77].

In language modeling, vector space word embeddings have been proposed in 2013 by Mikolov et al. [35,34]. They train neural networks for creating a low-dimensional, dense representation of words, which show two essential properties: (a) similar words are close in the vector space, and (b) relations between pairs of words can be represented as vectors as well, allowing for arithmetic operations in the vector space. In this work, we adapt those language modeling approaches for latent representation of entities in RDF graphs. Since language modeling techniques work on sentences, we first convert the graph into a set of sequences of entities using two different approaches, i.e., graph walks and Weisfeiler-Lehman Subtree RDF graph kernels. In the second step, we use those sequences to train a neural language model, which estimates the likelihood of a sequence of entities appearing in a graph. Once the training is finished, each entity in the graph is represented as a vector of latent numerical features. We show that the properties of word embeddings also hold for RDF entity embeddings, and that they can be exploited for various tasks.

We use several RDF graphs to show that such latent representation of entities have high relevance for different data mining and information retrieval tasks. The generation of the entities' vectors is task and dataset independent, i.e., we show that once the vectors are generated, they can be used for machine learning tasks, like classification and regression, entity and document modeling, and to estimate the closeness of items for content-based or hybrid recommender systems in a top- N scenario. Furthermore, since all entities are represented in a low dimensional feature space, building the learning models and algorithms becomes more efficient. To foster the reuse of the created feature sets, we provide the vector representations of DBpedia and Wikidata entities as ready-to-use files for download.

This paper considerably extends [61], in which we introduced RDF2Vec for the first time. In addition to the results presented in that paper, we show that the vector embeddings cannot only be used in machine learning tasks, but also for document modeling and recommender systems³, without a need to retrain the embedding models.

The rest of this paper is structured as follows. In Section 2, we give an overview of related work. In Section 3, we introduce our approach. In Section 4 through Section 7, we describe the evaluation setup and evaluate our approach on three different sets of tasks, i.e., machine learning, document modeling, and recommender systems. We conclude with a summary and an outlook on future work.

2. Related Work

While the representation of RDF as vectors in an embedding space itself is new, there is a larger body of related work in the three application areas discussed in this paper, i.e., the use of LOD in data mining, in document modeling, and in content-based recommender systems.

Generally, our work is closely related to the approaches DeepWalk [54] and Deep Graph Kernels [80]. DeepWalk uses language modeling approaches to learn social representations of vertices of graphs by modeling short random-walks on large social graphs, like BlogCatalog, Flickr, and YouTube. The Deep Graph Kernel approach extends the DeepWalk approach by modeling graph substructures, like graphlets, instead of random walks. The approach we propose in this paper differs from these two approaches in several aspects. First, we adapt the language modeling approaches on directed labeled RDF graphs, compared to the undirected graphs used in the approaches. Second, we show that task-independent entity vectors can be generated on large-scale knowledge graphs, which later can be reused on a variety of machine learning tasks on different datasets.

2.1. LOD in Machine Learning

In the recent past, a few approaches for generating data mining features from Linked Open Data have

³Preliminary results for the recommender task have already been published in [64], but this paper contains a deeper evaluation and discussion.

been proposed. Many of those approaches assume a manual design of the procedure for feature selection and, in most cases, this procedure results in the formulation of a SPARQL query by the user. LiDDM [24] allows the users to declare SPARQL queries for retrieving features from LOD that can be used in different machine learning techniques. Similarly, Cheng et al. [6] proposes an approach for automated feature generation after the user has specified the type of features, by specifying the SPARQL queries.

A similar approach has been used in the RapidMiner⁴ semweb plugin [25], which preprocesses RDF data in a way that can be further handled directly in RapidMiner. Mynarz et al. [40] have considered using user specified SPARQL queries in combination with SPARQL aggregates. *FeGeLOD* [52] and its successor, the *RapidMiner Linked Open Data Extension* [58], have been the first fully automatic unsupervised approach for enriching data with features that are derived from LOD. The approach uses six different unsupervised feature generation strategies, exploring specific or generic relations. It has been shown that such feature generation strategies can be used in many data mining tasks [53,58].

When dealing with *Kernel Functions* for graph-based data, we face similar problems as in feature generation and selection. Usually, the basic idea behind their computation is to evaluate the distance between two data instances by counting common substructures in the graphs of the instances, i.e. walks, paths and trees. In the past, many graph kernels have been proposed that are tailored towards specific applications [22,46], or towards specific semantic representations [12]. However, only a few approaches are general enough to be applied on any given RDF data, regardless the data mining task. Lösch et al. [32] introduce two general RDF graph kernels, based on intersection graphs and intersection trees. Later, the intersection tree path kernel was simplified by Vries et al. [76]. In another work, Vries et al. [75,77] introduce an approximation of the state-of-the-art Weisfeiler-Lehman graph kernel algorithm aimed at improving the computation time of the kernel when applied to RDF. Furthermore, the kernel implementation allows for explicit calculation of the instances' feature vectors, instead of pairwise similarities.

2.2. Entity and Document Modeling

Both for entity and document ranking, as well as for the subtask of computing the similarity or relatedness of entities and documents, different methods using LOD have been proposed.

2.2.1. Entity Relatedness

Semantic relatedness of entities has been heavily researched over the past couple of decades. There are two main direction of studies. The first are word-distributed-based approaches, which model entities as multi-dimensional vectors that are computed based on distributional semantics techniques [13,20,1]. The second are graph-based approaches relying on a graph structured KB, which are the focus of this paper.

Schumacher et al. [68] proposed one of the first approaches for entity ranking using the DBpedia knowledge graph. They use several path and graph based approaches for weighting the relations between entities, which are later used to calculate the entity relatedness. A similar approach is developed by Hulpus et al. [23], which uses local graph measures, targeted to the specific pair, while the previous approach uses global measures. More precisely, the authors propose the exclusivity-based relatedness measure that gives higher weights to relations that are less used in the graph. In [11] the authors propose a hybrid approach that exploits both textual and RDF data to rank resources in DBpedia related to the IT domain.

2.2.2. Entity and Document Similarity

As for the entity relatedness approaches, there are two main directions of research in the field of semantic document similarity, i.e., approaches based on word distributions, and graph-based approaches. Some of the earliest approaches of the first category make use of standard techniques like bag-of-words models, but also more sophisticated approaches. Explicit Semantic Analysis (ESA) [13] represents text as a vector of relevant concepts. Each concept corresponds to a Wikipedia article mapped into a vector space using the TF-IDF measure on the article's text. Similarly, Salient Semantic Analysis (SSA) [18] use hyperlinks within Wikipedia articles to other articles as vector features, instead of using the full body of text.

Nunes et al. [44] present a DBpedia based document similarity approach, in which they compute a document connectivity score based on document annotations, using measures from social network theory. Thiagarajan et al. [73] present a general framework showing how spreading activation can be used on seman-

⁴<http://www.rapidminer.com/>

tic networks to determine similarity of groups of entities. They experiment with Wordnet and Wikipedia Ontology as knowledge bases and determine similarity of generated user profiles based on a 1-1 annotation matching.

Schumacher et al. [68] use the same measure used for entity ranking (see above) to calculate semantic document similarity. Similarly, Paul et al. [48] present an approach for efficient semantic similarity computation that exploits hierarchical and transverse relations in the graph.

One approach that does not belong to these two main directions of research is the machine-learning approach by Huang et al. [21]. The approach proposes a measure that assesses similarity at both the lexical and semantic levels, and learns from human judgments how to combine them by using machine-learning techniques.

Our work is, to the best of our knowledge, the first to exploit the graph structure using neural language modeling for the purpose of entity relatedness and similarity.

2.3. Recommender Systems

Providing accurate suggestions, tailored to user’s needs and interests, is the main target of Recommender Systems (RS) [57], information filtering techniques commonly used to suggest items that are likely to be of use to a user. These techniques have proven to be very effective to face the *information overload* problem, that is the huge amount of information available on the Web, which risks to overwhelm user’s experience while retrieving items of interest. The numerous approaches facilitate the access to information in a personalized way, building a user profile and keeping it up-to-date.

RS address the information overload issue in two different ways, often combined into hybrid systems [4]: the *collaborative* approach [57] exploits information about the past behaviour and opinions of an existing user community to predict which item the current user will be more interested in, while the *content-based* approach [57] relies on the items “content”, that is the description of items’ characteristics. In a collaborative setting, a profile of the user is built by estimating her choice pattern through the behaviour of the overall user community. The content-based approach, instead, represents items by means of a set of features and defines a user profile as an assignment of importance to such features, exploiting the past interaction with the system. To overcome the limita-

tions of traditional approaches, which define the content based on partial metadata or on textual information optionally associated to an item, a process of “knowledge infusion” [69] has been performed for the last years, giving rise to the class of *semantics-aware* content-based recommender systems [15]. Many RS have incorporated ontological knowledge [33], unstructured or semi-structured knowledge sources (e.g., Wikipedia) [69], or the wealth of the LOD cloud, and recently the interest in unsupervised techniques where the human intervention is reduced or even withdrawn, has significantly increased.

LOD datasets, e.g., DBpedia [31], have been used in content-based recommender systems in [9] and [10]. The former performs a semantic expansion of the item content based on ontological information extracted from DBpedia and LinkedMDB [19], the first open semantic web database for movies, and tries to derive implicit relations between items. The latter involves both DBpedia and LinkedMDB and adapt the Vector Space Model to Linked Open Data: it represents the RDF graph as a 3-dimensional tensor where each slice is an ontological property (e.g. starring, director,...) and represents its adjacency matrix. It has been proved that leveraging LOD datasets is also effective for hybrid recommender systems [4], that is in those approaches that boost the collaborative information with additional knowledge, such as the item content. In [8] the authors propose *SPRank*, a hybrid recommendation algorithm that extracts semantic path-based features from DBpedia and uses them to compute top- N recommendations in a learning to rank approach and in multiple domains, movies, books and musical artists. *SPRank* is compared with numerous collaborative approaches based on matrix factorization [27,56] and with other hybrid RS, such as *BPR-SSLIM* [43], and exhibits good performance especially in those contexts characterized by high sparsity, where the contribution of the content becomes essential. Another hybrid approach is proposed in [59], which builds on training individual base recommenders and using global popularity scores as generic recommenders. The results of the individual recommenders are combined using stacking regression and rank aggregation. Most of these approaches can be referred to as *top-down* approaches [15], since they rely on the integration of external knowledge and cannot work without human intervention. On the other side, *bottom-up* approaches ground on the *distributional hypothesis* [17] for language modeling, according to which the meaning of words depends on the context in which they occur, in

some textual content. The resulting strategy is therefore unsupervised, requiring a corpora of textual documents for training as large as possible. Approaches based on the distributional hypothesis, referred to as *discriminative models*, behave as word embeddings techniques where each term (and document) becomes a point in the vector space. They substitute the term-document matrix typical of Vector Space Model with a term-context matrix on which they apply dimensionality reduction techniques such as Latent Semantic Indexing (LSI) [7] and the more scalable and incremental Random Indexing (RI) [65]. The latter has been involved in [38] and [39] to define the so called enhanced Vector Space Model (eVSM) for content-based RS, where user’s profile is incrementally built summing the features vectors representing documents liked by the user and a negation operator is introduced to take into account also negative preferences, inspired by [79], that is according to the principles of Quantum Logic.

Word embedding techniques are not limited to LSI and RI. The word2vec strategy has been recently presented in [34] and [35], and to the best of our knowledge, has been applied to item recommendations in a few works [37,47]. In particular, [37] is an empirical evaluation of LSI, RI and word2vec to make content-based movie recommendation exploiting textual information from Wikipedia, while [47] deals with check-in venue (location) recommendations and adds a non-textual feature, the past check-ins of the user. They both draw the conclusion that word2vec techniques are promising for the recommendation task. Finally, there is a single example of *product embedding* [16], namely *prod2vec*, which operates on the artificial graph of purchases, treating a purchase sequence as a “sentence” and products within the sequence as words.

3. Approach

In our approach, we adapt neural language models for RDF graph embeddings. Such approaches take advantage of the word order in text documents, explicitly modeling the assumption that closer words in a sequence are statistically more dependent. In the case of RDF graphs, we consider entities and relations between entities instead of word sequences. Thus, in order to apply such approaches on RDF graph data, we first have to transform the graph data into sequences of entities, which can be considered as sentences. Using those sentences, we can train the same neural language models to represent each entity in the RDF graph as a vector of numerical values in a latent feature space.

3.1. RDF Graph Sub-Structures Extraction

We propose two general approaches for converting graphs into a set of sequences of entities, i.e, graph walks and Weisfeiler-Lehman Subtree RDF Graph Kernels.

Definition 1 *An RDF graph is a graph $G = (V, E)$, where V is a set of vertices, and E is a set of directed edges.*

The objective of the conversion functions is for each vertex $v \in V$ to generate a set of sequences S_v , where the first token of each sequence $s \in S_v$ is the vertex v followed by a sequence of tokens, which might be edges, vertices, or any substructure extracted from the RDF graph, in an order that reflects the relations between the vertex v and the rest of the tokens, as well as among those tokens.

3.1.1. Graph Walks

In this approach, for a given graph $G = (V, E)$, for each vertex $v \in V$ we generate all graph walks P_v of depth d rooted in the vertex v . To generate the walks, we use the breadth-first algorithm. In the first iteration, the algorithm generates paths by exploring the direct outgoing edges of the root node v_r . The paths generated after the first iteration will have the following pattern $v_r \rightarrow e_{1i}$, where $i \in E(v_r)$. In the second iteration, for each of the previously explored edges the algorithm visits the connected vertices. The paths generated after the second iteration will follow the following pattern $v_r \rightarrow e_{1i} \rightarrow v_{1i}$. The algorithm continues until d iterations are reached. The final set of sequences for the given graph G is the union of the sequences of all the vertices $\bigcup_{v \in V} P_v$.

3.1.2. Weisfeiler-Lehman Subtree RDF Graph Kernels

In this approach, we use the subtree RDF adaptation of the Weisfeiler-Lehman algorithm presented in [75,77]. The Weisfeiler-Lehman Subtree graph kernel is a state-of-the-art, efficient kernel for graph comparison [70]. The kernel computes the number of subtrees shared between two (or more) graphs by using the Weisfeiler-Lehman test of graph isomorphism. This algorithm creates labels representing subtrees in h iterations.

There are two main modifications of the original Weisfeiler-Lehman graph kernel algorithm in order to be applicable on RDF graphs [77]. First, the RDF graphs have directed edges, which is reflected in the fact that the neighborhood of a vertex v contains only

the vertices reachable via outgoing edges. Second, as mentioned in the original algorithm, labels from two iterations can potentially be different while still representing the same subtree. To make sure that this does not happen, the authors in [77] have added tracking of the neighboring labels in the previous iteration, via the multiset of the previous iteration. If the multiset of the current iteration is identical to that of the previous iteration, the label of the previous iteration is reused.

The procedure of converting the RDF graph to a set of sequences of tokens goes as follows: (i) for a given graph $G = (V, E)$, we define the Weisfeiler-Lehman algorithm parameters, i.e., the number of iterations h and the vertex subgraph depth d , which defines the subgraph in which the subtrees will be counted for the given vertex; (ii) after each iteration, for each vertex $v \in V$ of the original graph G , we extract all the paths of depth d within the subgraph of the vertex v on the relabeled graph. We set the original label of the vertex v as the starting token of each path, which is then considered as a sequence of tokens. The sequences after the first iteration will have the following pattern $v_r \rightarrow T_1 \rightarrow T_1 \dots T_d$, where T_d is a subtree that appears on depth d in the vertex's subgraph; (iii) we repeat step (ii) until the maximum iterations h are reached. (iv) The final set of sequences is the union of the sequences of all the vertices in each iteration $\bigcup_{i=1}^h \bigcup_{v \in V} P_v$.

3.2. Neural Language Models – word2vec

Neural language models have been developed in the NLP field as an alternative to represent texts as a bag of words, and hence, a binary feature vector, where each vector index represents one word. While such approaches are simple and robust, they suffer from several drawbacks, e.g., high dimensionality and severe data sparsity, which limits the performances of such techniques. To overcome such limitations, neural language models have been proposed, inducing low-dimensional, distributed embeddings of words by means of neural networks. The goal of such approaches is to estimate the likelihood of a specific sequence of words appearing in a corpus, explicitly modeling the assumption that closer words in the word sequence are statistically more dependent.

While some of the initially proposed approaches suffered from inefficient training of the neural network models, with the recent advancements in the field several efficient approaches has been proposed. One of the most popular and widely used is the word2vec neural language model [34,35]. Word2vec is a particularly

computationally-efficient two-layer neural net model for learning word embeddings from raw text. There are two different algorithms, the Continuous Bag-of-Words model (CBOW) and the Skip-Gram model.

3.2.1. Continuous Bag-of-Words Model

The CBOW model predicts target words from context words within a given window. The model architecture is shown in Fig. 1a. The input layer is comprised from all the surrounding words for which the input vectors are retrieved from the input weight matrix, averaged, and projected in the projection layer. Then, using the weights from the output weight matrix, a score for each word in the vocabulary is computed, which is the probability of the word being a target word. Formally, given a sequence of training words $w_1, w_2, w_3, \dots, w_T$, and a context window c , the objective of the CBOW model is to maximize the average log probability:

$$\frac{1}{T} \sum_{t=1}^T \log p(w_t | w_{t-c} \dots w_{t+c}), \quad (1)$$

where the probability $p(w_t | w_{t-c} \dots w_{t+c})$ is calculated using the softmax function:

$$p(w_t | w_{t-c} \dots w_{t+c}) = \frac{\exp(\bar{v}^T v'_w)}{\sum_{w=1}^V \exp(\bar{v}^T v'_w)}, \quad (2)$$

where v'_w is the output vector of the word w , V is the complete vocabulary of words, and \bar{v} is the averaged input vector of all the context words:

$$\bar{v} = \frac{1}{2c} \sum_{-c \leq j \leq c, j \neq 0} v_{w_{t+j}} \quad (3)$$

3.2.2. Skip-Gram Model

The skip-gram model does the inverse of the CBOW model and tries to predict the context words from the target words (Fig. 1b). More formally, given a sequence of training words $w_1, w_2, w_3, \dots, w_T$, and a context window c , the objective of the skip-gram model is to maximize the following average log probability:

$$\frac{1}{T} \sum_{t=1}^T \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t), \quad (4)$$

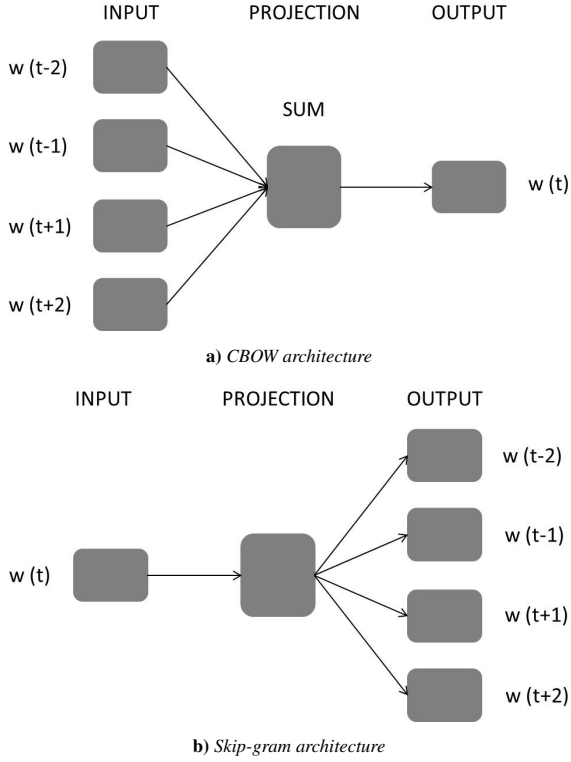


Fig. 1. Architecture of the CBOW and Skip-gram model.

where the probability $p(w_{t+j}|w_t)$ is calculated using the softmax function:

$$p(w_{t+c}|w_t) = \frac{\exp(v_{w_{t+c}}^T v_{w_t})}{\sum_{v=1}^V \exp(v_v^T v_{w_t})}, \quad (5)$$

where v_w and v'_w are the input and the output vector of the word w , and V is the complete vocabulary of words.

In both cases, calculating the softmax function is computationally inefficient, as the cost for computing is proportional to the size of the vocabulary. Therefore, two optimization techniques have been proposed, i.e., hierarchical softmax and negative sampling [35]. Empirical studies have shown that in most cases negative sampling leads to a better performance than hierarchical softmax, which depends on the selected negative samples, but it has higher runtime.

Once the training is finished, all words (or, in our case, entities) are projected into a lower-dimensional feature space, and semantically similar words (or entities) are positioned close to each other.

4. Evaluation

We evaluate our approach on three different tasks: (i) standard machine-learning classification and regression; (ii) document similarity and entity relatedness; (iii) top- N recommendation task both with content-based and hybrid RSs. For all three tasks, we utilize two of the most prominent RDF knowledge graphs [50], i.e., DBpedia [31] and Wikidata [74]. DBpedia is a knowledge graph which is extracted from structured data in Wikipedia. The main source for this extraction are the key-value pairs in the Wikipedia infoboxes. Wikidata is a collaboratively edited knowledge graph, operated by the Wikimedia foundation⁵ that also hosts various language editions of Wikipedia.

We use the English version of the 2015-10 DBpedia dataset, which contains 4,641,890 instances and 1,369 mapping-based properties. In our evaluation we only consider object properties, and ignore datatype properties and literals.

For the Wikidata dataset we use the simplified and derived RDF dumps from 2016-03-28⁶. The dataset contains 17,340,659 entities in total. As for the DBpedia dataset, we only consider object properties, and ignore the data properties and literals.

The first step of our approach is to convert the RDF graphs into a set of sequences. As the number of generated walks increases exponentially [77] with the graph traversal depth, calculating Weisfeiler-Lehman subtrees RDF kernels, or all graph walks with a given depth d for all of the entities in the large RDF graph quickly becomes unmanageable. Therefore, to extract the entities embeddings for the large RDF datasets, we use only random graph walks entity sequences. More precisely, we follow the approach presented in [54] to generate limited number of random walks for each entity. For DBpedia, we experiment with 500 walks per entity with depth of 4 and 8, while for Wikidata, we use only 200 walks per entity with depth of 4. Additionally, for each entity in DBpedia and Wikidata, we include all the walks of depth 2, i.e., direct outgoing relations. We use the corpora of sequences to build both CBOW and Skip-Gram models with the following parameters: window size = 5; number of iterations = 5; negative sampling for optimization; negative samples = 25; with average input vector for CBOW. We exper-

⁵<http://wikimediafoundation.org/>

⁶http://tools.wmflabs.org/wikidata-exports/rdf/index.php?content=dump__download.php\&dump=20160328

iment with 200 and 500 dimensions for the entities’ vectors. All the models, as well as the code, are publicly available.⁷

5. Machine Learning with Background Knowledge from LOD

Linking entities in a machine learning task to those in the LOD cloud helps generating additional features, which may help improving the overall learning outcome. For example, when learning a predictive model for the success of a movie, adding knowledge from the LOD cloud (such as the movie’s budget, director, genre, etc.) can lead to a more accurate model.

5.1. Experimental Setup

For evaluating the performance of our RDF embeddings in machine learning tasks, we perform an evaluation on a set of benchmark datasets. The dataset contains of three smaller-scale RDF datasets (i.e., AIFB, MUTAG, and BGS), where the classification target is the value of a selected property within the dataset, and five larger datasets linked to DBpedia and Wikidata, where the target is an external variable (e.g., the metacritic score of an album or a movie). The latter datasets are used both for classification and regression. Details on the datasets can be found in [63].

For each of the small RDF datasets, we first build two corpora of sequences, i.e., the set of sequences generated from graph walks with depth 8 (marked as W2V), and set of sequences generated from Weisfeiler-Lehman subtree kernels (marked as K2V). For the Weisfeiler-Lehman algorithm, we use 4 iterations and depth of 2, and after each iteration we extract all walks for each entity with the same depth. We use the corpora of sequences to build both CBOW and Skip-Gram models with the following parameters: window size = 5; number of iterations = 10; negative sampling for optimization; negative samples = 25; with average input vector for CBOW. We experiment with 200 and 500 dimensions for the entities’ vectors.

We use the RDF embeddings of DBpedia and Wikidata (see Section 4) on the five larger datasets, which provide classification/regression targets for DBpedia/Wikidata entities (see Table 1).

⁷<http://data.dws.informatik.uni-mannheim.de/rdf2vec/>

Table 1

Datasets overview. For each dataset, we depict the number of instances, the machine learning tasks in which the dataset is used (*C* stands for classification, and *R* stands for regression) and the source of the dataset.

Dataset	#Instances	ML Task	Original Source
Cities	212	R/C (c=3)	Mercer
Metacritic Albums	1600	R/C (c=2)	Metacritic
Metacritic Movies	2000	R/C (c=2)	Metacritic
AAUP	960	R/C (c=3)	JSE
Forbes	1585	R/C (c=3)	Forbes
AIFB	176	C (c=4)	AIFB
MUTAG	340	C (c=2)	MUTAG
BGS	146	C (c=2)	BGS

We compare our approach to several baselines. For generating the data mining features, we use three strategies that take into account the direct relations to other resources in the graph [52], and two strategies for features derived from graph sub-structures [77]:

- Features derived from specific relations. In the experiments we use the relations *rdf:type* (types), and *dcterms:subject* (categories) for datasets linked to DBpedia.
- Features derived from generic relations, i.e., we generate a feature for each incoming (rel in) or outgoing relation (rel out) of an entity, ignoring the value or target entity of the relation.
- Features derived from generic relations-values, i.e, we generate feature for each incoming (rel-vals in) or outgoing relation (rel-vals out) of an entity including the value of the relation.
- Kernels that count substructures in the RDF graph around the instance node. These substructures are explicitly generated and represented as sparse feature vectors.
 - * The Weisfeiler-Lehman (WL) graph kernel for RDF [77] counts full subtrees in the subgraph around the instance node. This kernel has two parameters, the subgraph depth d and the number of iterations h (which determines the depth of the subtrees). We use two pairs of settings, $d = 1, h = 2$ and $d = 2, h = 3$.
 - * The Intersection Tree Path kernel for RDF [77] counts the walks in the subtree that spans from the instance node. Only the walks that go through the instance node are considered. We will therefore refer to it as the root Walk Count (WC) kernel. The root WC kernel has one parameter: the length of the paths l , for which we test 2 and 3.

We perform two learning tasks, i.e., classification and regression. For classification tasks, we use Naive Bayes, k-Nearest Neighbors ($k=3$), C4.5 decision tree, and Support Vector Machines (SVMs). For the SVM classifier we optimize the parameter C in the range $\{10^{-3}, 10^{-2}, 0.1, 1, 10, 10^2, 10^3\}$. For regression, we use Linear Regression, M5Rules, and k-Nearest Neighbors ($k=3$). We measure accuracy for classification tasks, and root mean squared error (RMSE) for regression tasks. The results are calculated using stratified 10-fold cross validation.

The strategies for creating propositional features from Linked Open Data are implemented in the RapidMiner LOD extension⁸ [53,58]. The experiments, including the feature generation and the evaluation, were performed using the RapidMiner data analytics platform.⁹ The RapidMiner processes and the complete results can be found online.¹⁰

5.2. Results

The results for the task of classification on the small RDF datasets are shown in Table 2. From the results we can observe that the K2V approach outperforms all the other approaches. More precisely, using the skip-gram feature vectors of size 500 in an SVM model provides the best results on all three datasets. The W2V approach on all three datasets performs closely to the standard graph substructure feature generation strategies, but it does not outperform them. K2V outperforms W2V because it is able to capture more complex substructures in the graph, like sub-trees, while W2V focuses only on graph paths.

The results for the task of classification on the five datasets using the DBpedia and Wikidata entities' vectors are shown in Table 3, and the results for the task of regression on the five datasets using the DBpedia and Wikidata entities' vectors are given in Table 4. We can observe that the latent vectors extracted from DBpedia and Wikidata outperform all of the standard feature generation approaches. In general, the DBpedia vectors work better than the Wikidata vectors, where the skip-gram vectors with size 200 or 500 built on graph walks of depth 8 on most of the datasets lead to the best performances. An exception is the AAUP dataset,

where the Wikidata skip-gram 500 vectors outperform the other approaches.

On both tasks, we can observe that the skip-gram vectors perform better than the CBOW vectors. Furthermore, the vectors with higher dimensionality and longer paths lead to a better representation of the entities and better performance on most of the datasets. However, for the variety of tasks at hand, there is no universal approach, i.e., embedding model and a machine learning method, that consistently outperforms the others.

5.3. Semantics of Vector Representations

To analyse the semantics of the vector representations, we employ Principal Component Analysis (PCA) to project the entities' feature vectors into a two dimensional feature space. We selected seven countries and their capital cities, and visualized their vectors as shown in Figure 4. Figure 2a shows the corresponding DBpedia vectors, and Figure 2b shows the corresponding Wikidata vectors. The figure illustrates the ability of the model to automatically organize entities of different types, and preserve the relationships between different entities. For example, we can see that there is a clear separation between the countries and the cities, and the relation "capital" between each pair of country and the corresponding capital city is preserved. Furthermore, we can observe that more similar entities are positioned closer to each other, e.g., we can see that the countries that are part of the EU are closer to each other, and the same applies for the Asian countries.

5.4. Features Increase Rate

Finally, we conduct a scalability experiment, where we examine how the number of instances affects the number of generated features by each feature generation strategy. For this purpose we use the *Metacritic Movies* dataset. We start with a random sample of 100 instances, and in each next step we add 200 (or 300) unused instances, until the complete dataset is used, i.e., 2,000 instances. The number of generated features for each sub-sample of the dataset using each of the feature generation strategies is shown in Figure 3.

From the chart, we can observe that the number of generated features sharply increases when adding more samples in the datasets, especially for the strategies based on graph substructures. However, the number of features remains the same when using the

⁸<http://dws.informatik.uni-mannheim.de/en/research/rapidminer-lod-extension>

⁹<https://rapidminer.com/>

¹⁰http://data.dws.informatik.uni-mannheim.de/rmlod/LOD_ML_Datasets/

Table 2

Classification results on the small RDF datasets. The best results are marked in bold. Experiments marked with “\” did not finish within ten days, or have run out of memory.

Strategy/Dataset	AIFB				MUTAG				BGS			
	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5
rel in	16.99	47.19	50.70	50.62	\	\	\	\	61.76	54.67	63.76	63.76
rel out	45.07	45.56	50.70	51.76	41.18	54.41	62.94	62.06	54.76	69.05	72.70	69.33
rel in & out	25.59	51.24	50.80	51.80	\	\	\	\	54.76	67.00	72.00	70.00
rel-vals in	73.24	54.54	81.86	80.75	\	\	\	\	79.48	83.52	86.50	68.57
rel-vals out	86.86	55.69	82.39	71.73	62.35	62.06	73.53	62.94	84.95	65.29	83.10	73.38
rel-vals in&out	87.42	57.91	88.57	85.82	\	\	\	\	84.95	70.81	85.80	72.67
WL_1_2	85.69	53.30	92.68	71.08	91.12	62.06	92.59	93.29	85.48	63.62	82.14	75.29
WL_2_2	85.65	65.95	83.43	89.25	70.59	62.06	94.29	93.47	90.33	85.57	91.05	87.67
WC_2	86.24	60.27	75.03	71.05	90.94	62.06	91.76	93.82	84.81	69.00	83.57	76.90
WC_3	86.83	64.18	82.97	71.05	92.00	72.56	86.47	93.82	85.00	67.00	78.71	76.90
W2V CBOW 200	70.00	69.97	79.48	65.33	74.71	72.35	80.29	74.41	56.14	74.00	74.71	67.38
W2V CBOW 500	69.97	69.44	82.88	73.40	75.59	70.59	82.06	72.06	55.43	73.95	74.05	65.86
W2V SG 200	76.76	71.67	87.39	65.36	70.00	71.76	77.94	68.53	66.95	69.10	75.29	71.24
W2V SG 500	76.67	76.18	89.55	71.05	72.35	72.65	78.24	68.24	68.38	71.19	78.10	63.00
K2V CBOW 200	85.16	84.48	87.48	76.08	78.82	69.41	86.47	68.53	93.14	95.57	94.71	88.19
K2V CBOW 500	90.98	88.17	86.83	76.18	80.59	70.88	90.88	66.76	93.48	95.67	94.82	87.26
K2V SG 200	85.65	87.96	90.82	75.26	78.53	69.29	95.88	66.00	91.19	93.24	95.95	87.05
K2V SG 500	88.73	88.66	93.41	69.90	82.06	70.29	96.18	66.18	91.81	93.19	96.33	80.76




Fig. 3. Features increase rate per strategy (log scale).

RDF2Vec approach, independently of the number of samples in the data. Thus, by design, it scales to larger datasets without increasing the dimensionality of the dataset.

6. Entity and Document Modeling

Calculating entity relatedness and similarity are fundamental problems in numerous tasks in information


retrieval, natural language processing, and Web-based knowledge extractions. While similarity only considers subsumption relations to assess how two objects are alike, relatedness takes into account a broader range of relations, i.e., the notion of relatedness is wider than that of similarity. For example, “Facebook” and “Google” are both entities of the class company, and they have high similarity and relatedness score. On the other hand, “Facebook” and “Mark Zuckerberg” are not similar at all, but are highly related. While

Table 3: Classification results. The first number represents the dimensionality of the vectors, while the second number represent the value for the depth parameter. The best results are marked in bold. Experiments marked with “\” did not finish within ten days, or have run out of memory.


Strategy/Dataset	Cities				Metaelectric Movies				Metaelectric Albums				AAUP				Forbes			
	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5	NB	KNN	SVM	C4.5
types	55.71	56.17	63.21	59.05	68.00	57.60	71.40	70.00	66.50	50.75	62.31	54.44	41.00	85.62	91.67	92.78	55.08	75.84	75.67	75.85
categories	55.74	49.98	62.39	56.17	75.25	62.70	76.35	69.50	67.40	54.13	64.50	56.62	48.00	85.83	90.78	91.87	60.38	76.11	75.70	75.70
rel in	60.41	58.46	71.70	60.35	52.75	49.90	60.35	60.10	51.13	62.19	65.25	60.75	45.63	85.94	90.62	92.81	50.24	76.49	75.16	76.10
rel out	47.62	60.00	66.04	56.71	52.90	58.45	66.40	62.70	58.75	63.75	62.25	64.50	41.15	85.83	89.58	91.35	64.73	75.84	75.73	75.92
rel in & out	59.44	58.57	66.04	56.47	52.95	59.30	67.75	62.55	58.69	64.50	67.38	61.56	42.71	85.94	89.67	92.50	22.27	75.96	76.34	75.98
rel-vals in	\	\	\	\	50.60	50.00	50.60	50.00	50.88	50.00	50.81	50.00	54.06	84.69	89.51	\	14.95	76.15	76.97	75.73
rel-vals in&out	53.79	35.91	55.66	64.13	78.50	54.78	78.71	\	74.06	52.56	76.99	\	57.81	85.73	91.46	91.78	67.09	75.61	75.74	76.74
WL_1_2	70.98	49.31	65.34	75.29	75.45	66.90	79.30	70.80	73.63	64.69	76.25	62.00	58.33	91.04	91.46	92.40	64.17	75.71	75.10	76.59
WL_2_3	65.48	53.29	69.90	69.31	\	\	\	\	\	\	\	\	\	\	\	\	\	\	\	\
WC_2	72.71	47.39	66.48	75.13	75.39	65.89	74.93	69.08	72.00	60.63	76.88	63.69	57.29	90.63	93.44	92.60	64.23	75.77	76.22	76.47
WC_3	65.52	52.36	67.95	65.15	74.25	55.30	78.40	\	72.81	52.87	77.94	\	57.19	90.73	90.94	92.60	64.04	75.65	76.22	76.59
DB2vec CBOW 200 4	59.32	68.84	77.39	64.32	65.60	79.74	82.90	74.33	70.72	71.86	76.36	67.24	73.36	89.65	29.00	92.45	89.38	80.94	76.83	84.81
DB2vec CBOW 500 4	59.32	71.34	76.37	66.34	65.65	79.49	82.75	73.87	69.71	71.93	75.41	65.65	72.71	89.65	29.11	92.01	89.02	80.82	76.95	85.17
DB2vec SG 200 4	60.34	71.82	76.37	65.37	65.25	80.44	83.25	73.87	68.95	73.89	76.11	67.87	71.20	89.65	28.90	92.12	88.78	80.82	77.92	85.77
DB2vec SG 500 4	58.34	72.84	76.87	67.84	65.45	80.14	83.65	72.82	70.41	74.34	78.44	67.49	71.19	89.65	28.90	92.23	88.30	80.94	77.25	84.81
DB2vec CBOW 200 8	69.26	69.87	67.32	63.13	57.83	70.08	65.25	67.47	67.91	64.44	72.42	65.39	68.18	85.33	28.90	90.50	77.35	80.34	28.90	85.17
DB2vec CBOW 500 8	62.26	69.87	76.84	63.21	58.78	69.82	69.46	67.67	67.53	65.83	74.26	63.42	62.90	85.22	29.11	90.61	89.86	80.34	78.65	84.81
DB2vec SG 200 8	73.32	75.89	78.92	60.74	79.94	79.49	83.30	75.13	77.25	76.87	79.72	69.14	78.53	85.12	29.22	91.04	90.10	80.38	78.96	84.68
DB2vec SG 500 8	89.73	69.16	84.19	72.25	80.24	78.68	82.80	72.42	73.57	76.30	78.20	68.70	75.07	94.48	29.11	94.15	88.53	80.58	77.79	86.38
WD2vec CBOW 200 4	68.76	57.71	75.56	61.37	51.49	52.20	51.64	49.01	50.86	50.29	51.44	50.09	50.54	90.18	89.63	88.83	49.84	81.08	76.77	79.14
WD2vec CBOW 500 4	68.24	57.75	85.56	64.54	49.22	48.56	51.04	50.98	53.08	50.03	52.33	53.28	48.45	90.39	89.74	88.31	51.95	80.74	77.18	80.32
WD2vec SG 200 4	72.58	57.53	75.48	52.32	69.53	70.14	75.39	67.00	60.32	62.03	64.76	58.54	60.87	90.50	89.63	89.98	65.45	81.17	77.74	77.03
WD2vec SG 300 4	83.20	60.72	79.87	61.67	71.10	70.19	76.30	67.31	55.31	58.92	63.42	56.63	55.85	90.60	89.63	87.69	58.95	81.17	79.00	79.56

Table 4: Regression results. The first number represents the dimensionality of the vectors, while the second number represent the value for the depth parameter. The best results are marked in bold. Experiments that did not finish within ten days, or that have run out of memory are marked with “\”.

Strategy/Dataset	Cities					Metacritic Movies					Metacritic Albums					AAUP					Forbes				
	LR	KNN	M5	LR	KNN	M5	LR	KNN	M5	LR	KNN	M5	LR	KNN	M5	LR	KNN	M5	LR	KNN	M5				
types	24.30	22.16	18.79	77.80	30.68	22.16	16.45	18.36	13.95	9.83	34.95	6.28	29.22	21.07	18.32										
categories	18.88	22.68	22.32	84.57	23.87	22.50	16.73	16.64	13.95	8.08	34.94	6.16	19.16	21.48	18.39										
rel in	49.87	18.53	19.21	22.60	41.40	22.56	13.50	22.06	13.43	9.69	34.98	6.56	27.56	20.93	18.60										
rel out	49.87	18.53	19.21	21.45	24.42	20.74	13.32	14.59	13.06	8.82	34.95	6.32	21.73	21.11	18.97										
rel in & out	40.80	18.21	18.80	21.45	24.42	20.74	13.33	14.52	12.91	12.97	34.95	6.36	26.44	20.98	19.54										
rel-vals in	\	\	\	21.46	24.19	20.43	13.94	23.05	13.95	\	34.96	6.27	\	20.86	19.31										
rel-vals out	20.93	23.87	20.97	25.99	32.18	22.93	\	15.28	13.34	\	34.95	6.18	\	20.48	18.37										
rel-vals in&out	\	\	\	\	25.37	20.96	\	15.47	13.33	\	34.94	6.18	\	20.20	18.20										
WL_1_2	20.21	24.60	20.85	\	21.62	19.84	\	13.99	12.81	\	34.96	6.27	\	19.81	19.49										
WL_2_3	17.79	20.42	17.04	\	\	\	\	\	\	\	\	\	\	\	\										
WC_2	20.33	25.95	19.55	\	22.80	22.99	\	14.54	12.87	9.12	34.95	6.24	\	20.45	19.26										
WC_3	19.51	33.16	19.05	\	23.86	19.19	\	19.51	13.02	\	35.39	6.31	\	20.58	19.04										
DB2vec CBOW 200 4	14.37	12.55	14.33	15.90	17.46	15.89	11.79	12.45	11.59	12.13	45.76	12.00	18.32	26.19	17.43										
DB2vec CBOW 500 4	14.99	12.46	14.66	15.90	17.45	15.73	11.49	12.60	11.48	12.44	45.67	12.30	18.23	26.27	17.62										
DB2vec SG 200 4	13.38	12.54	15.13	15.81	17.07	15.84	11.30	12.36	11.42	12.13	45.72	12.10	17.63	26.13	17.85										
DB2vec SG 500 4	14.73	13.25	16.80	15.66	17.14	15.67	11.20	12.11	11.28	12.09	45.76	11.93	18.23	26.09	17.74										
DB2vec CBOW 200 8	16.17	17.14	17.56	21.55	23.75	21.46	13.35	15.41	13.43	6.47	55.76	6.47	24.17	26.48	22.61										
DB2vec CBOW 500 8	18.13	17.19	18.50	20.77	23.67	20.69	13.20	15.14	13.25	6.54	55.33	6.55	21.16	25.90	20.33										
DB2vec SG 200 8	12.85	14.95	12.92	15.15	17.13	15.12	10.90	11.43	10.90	6.22	56.95	6.25	18.66	21.20	18.57										
DB2vec SG 500 8	11.92	12.67	10.19	15.45	17.80	15.50	10.89	11.72	10.97	6.26	56.95	6.29	18.35	21.04	16.61										
WD2vec CBOW 200 4	20.15	17.52	20.02	23.54	25.90	23.39	14.73	16.12	14.55	16.80	42.61	6.60	27.48	22.60	21.77										
WD2vec CBOW 500 4	23.76	18.33	20.39	24.14	22.18	24.56	14.09	16.09	14.00	13.08	42.89	6.08	50.23	21.92	26.66										
WD2vec SG 200 4	20.47	18.69	20.72	19.72	21.44	19.10	13.51	13.91	13.67	6.86	42.82	6.52	23.69	21.59	20.49										
WD2vec SG 500 4	22.25	19.41	19.23	25.99	21.26	19.19	13.23	14.96	13.25	8.27	42.84	6.05	21.98	21.73	21.58										



a) DBpedia vectors



b) Wikidata vectors

Fig. 2. Two-dimensional PCA projection of the 500-dimensional Skip-gram vectors of countries and their capital cities.

“Google” and “Mark Zuckerberg” are not similar at all, and have low relatedness value.

6.1. Approach

In this section we introduce several approaches for entity and document modeling based on the previously built entity latent vectors.

6.1.1. Entity Similarity

As previously mentioned, in the RDF2vec feature embedding space (see Section 3) semantically similar

entities appear close to each other in the feature space. Therefore, the problem of calculating the similarity between two instances is a matter of calculating the distance between two instances in the given feature space. To do so, we use the standard cosine similarity measure, which is applied on the vectors of the entities. Formally, the similarity between two entities e_1 and e_2 , with vectors V_1 and V_2 , is calculated as the cosine similarity between the vectors V_1 and V_2 :

$$\text{sim}(e_1, e_2) = \frac{V_1 \cdot V_2}{\|V_1\| \cdot \|V_2\|} \quad (6)$$

6.1.2. Document Similarity

We use the entity similarity approach in the task of calculating semantic document similarity. We follow similar approach as the one presented in [48], where two documents are considered to be similar if many entities of the one document are similar to at least one entity in the other document. More precisely, we try to identify the most similar pairs of entities in both documents, ignoring the similarity of all the other 1-1 similarities values.

Given two documents d_1 and d_2 , the similarity between the documents $\text{sim}(d_1, d_2)$ is calculated as follows:

1. Extract the sets of entities E_1 and E_2 in the documents d_1 and d_2 .
2. Calculate the similarity score $\text{sim}(e_{1i}, e_{2j})$ for each pair of entities in document d_1 and d_2 , where $e_{1i} \in E_1$ and $e_{2j} \in E_2$.
3. For each entity e_{1i} in d_1 identify the maximum similarity to an entity in d_2 $\text{max_sim}(e_{1i}, e_{2j} \in E_2)$, and vice versa.
4. Calculate the similarity score between the documents d_1 and d_2 as:

$$\text{sim}(d_1, d_2) = \frac{\sum_{i=1}^{|E_1|} \text{max_sim}(e_{1i}, e_{2j} \in E_2) + \sum_{j=1}^{|E_2|} \text{max_sim}(e_{2j}, e_{1i} \in E_1)}{|E_1| + |E_2|} \quad (7)$$

6.1.3. Entity Relatedness

In this approach we assume that two entities are related if they often appear in the same context. For example, “Facebook” and “Mark Zuckerberg”, which are highly related, are often used in the same context in many sentences. To calculate the probability of two entities being in the same context, we make use of the RDF2Vec models and the set of sequences of entities generated as described in Section 3. Given a RDF2vec

model and a set of sequences of entities, we calculate the relatedness between two entities e_1 and e_2 , as the probability $p(e_1|e_2)$ calculated using the softmax function. In the case of a CBOW model, the probability is calculated as:

$$p(e_1|e_2) = \frac{\exp(v_{e_2}^T v'_{e_1})}{\sum_{e=1}^V \exp(v_{e_2}^T v'_e)},$$

where v'_e is the output vector of the entity e , and V is the complete vocabulary of entities.

In the case of a skip-gram model, the probability is calculated as:

$$p(e_1|e_2) = \frac{\exp(v_{e_1}^T v_{e_2})}{\sum_{e=1}^V \exp(v_e^T v_{e_2})}, \quad (8)$$

where v_e and v'_e are the input and the output vector of the entity e , and V is the complete vocabulary of entities.

6.2. Evaluation

For both tasks of determining entity relatedness and document similarity, benchmark datasets exist. We use those datasets to compare the use of RDF2Vec models against state of the art approaches.

6.2.1. Entity Relatedness

For evaluating the entity relatedness approach, we use the KORE dataset [20]. The dataset consists of 21 main entities, whose relatedness to the other 20 entities each has been manually assessed, leading to 420 rated entity pairs. We use the Spearman’s rank correlation as an evaluation metric.

We use two approaches for calculating the relatedness rank between the entities, i.e. (i) the entity similarity approach described in section 6.1.1; (ii) the entity relatedness approach described in section 6.1.3.

We evaluate each of the RDF2Vec models separately. Furthermore, we also compare to the Wiki2vec model¹¹, which is built on the complete Wikipedia corpus, and provides vectors for each DBpedia entity.

Table 5 shows the Spearman’s rank correlation results when using the entity similarity approach. Table 6 shows the results for the relatedness approach. The results show that the skip-gram models built on DBpedia with eight hops lead to the best performances,

where the relatedness approach outperforms the similarity approach.

Furthermore, we compare our approaches to several state-of-the-art graph-based entity relatedness approaches:

- baseline: computes entity relatedness as a function of distance between the entities in the network.
- KORE: calculates keyphrase overlap relatedness, as described in the original KORE paper [20].
- CombIC: semantic similarity using a Graph Edit Distance based measure [68].
- ER: Exclusivity-based relatedness [23].

The comparison shows that our entity relatedness approach outperforms all the rest for each category of entities. Interestingly enough, the entity *similarity* approach, although addressing a different task, also outperforms the majority of state of the art approaches.

6.3. Document Similarity

To evaluate the document similarity approach, we use the LP50 dataset [29], namely a collection of 50 news articles from the Australian Broadcasting Corporation (ABC), which were pairwise annotated with similarity rating on a Likert scale from 1 (very different) to 5 (very similar) by 8 to 12 different human annotators. To obtain the final similarity judgments, the scores are averaged for each pair the scores of all annotators. As a evaluation metrics we use Pearson’s linear correlation coefficient and Spearman’s rank correlation plus their harmonic mean.

Again, we first evaluate each of the RDF2Vec models separately. Table 8 shows document similarity results. As for the entity relatedness, the results show that the skip-gram models built on DBpedia with 8 hops lead to the best performances.

Furthermore, we compare our approach to several state-of-the-art graph-based document similarity approaches:

- TF-IDF: Distributional baseline algorithm.
- AnnOv: Similarity score based on annotation overlap.
- Explicit Semantic Analysis (ESA) [13].
- GED: semantic similarity using a Graph Edit Distance based measure [68].
- Salient Semantic Analysis (SSA), Latent Semantic Analysis (LSA) [18].
- Graph-based Semantic Similarity (GBSS) [48].

¹¹<https://github.com/idio/wiki2vec>

Table 5
Similarity-based relatedness Spearman's rank correlation results

Model	IT companies	Hollywood Celebrities	Television Series	Video Games	Chuck Norris	All 21 entities
DB2vec CBOW 200 4	0.538	0.560	0.572	0.596	0.500	0.564
DB2vec CBOW 500 4	0.546	0.544	0.564	0.606	0.496	0.562
DB2vec SG 200 4	0.508	0.546	0.497	0.634	0.570	0.547
DB2vec SG 500 4	0.507	0.538	0.505	0.611	0.588	0.542
DB2vec CBOW 200 8	0.611	0.495	0.315	0.443	0.365	0.461
DB2vec CBOW 500 8	0.486	0.507	0.285	0.440	0.470	0.432
DB2vec SG 200 8	0.739	0.723	0.526	0.659	0.625	0.660
DB2vec SG 500 8	0.743	0.734	0.635	0.669	0.628	0.692
WD2vec CBOW 200 4	0.246	0.418	0.156	0.374	0.409	0.304
WD2vec CBOW 500 4	0.190	0.403	0.103	0.106	0.150	0.198
WD2vec SG 200 4	0.502	0.604	0.405	0.578	0.279	0.510
WD2vec SG 500 4	0.464	0.562	0.313	0.465	0.168	0.437
Wiki2vec	0.613	0.544	0.334	0.618	0.436	0.523

Table 6
Context-based relatedness Spearman's rank correlation results

Model	IT companies	Hollywood Celebrities	Television Series	Video Games	Chuck Norris	All 21 entities
DB2vec CBOW 200 4	0.671	0.566	0.591	0.434	0.609	0.568
DB2vec CBOW 500 4	0.672	0.622	0.578	0.440	0.581	0.578
DB2vec SG 200 4	0.666	0.449	0.611	0.360	0.630	0.526
DB2vec SG 500 4	0.667	0.444	0.609	0.389	0.668	0.534
DB2vec CBOW 200 8	0.579	0.484	0.368	0.460	0.412	0.470
DB2vec CBOW 500 8	0.552	0.522	0.302	0.487	0.665	0.475
DB2vec SG 200 8	0.811	0.778	0.711	0.658	0.670	0.736
DB2vec SG 500 8	0.748	0.729	0.689	0.537	0.625	0.673
WD2vec CBOW 200 4	0.287	0.241	-0.025	0.311	0.226	0.205
WD2vec CBOW 500 4	0.166	0.215	0.233	0.335	0.344	0.243
WD2vec SG 200 4	0.574	0.671	0.504	0.410	0.079	0.518
WD2vec SG 500 4	0.661	0.639	0.537	0.395	0.474	0.554
Wiki2vec	0.291	0.296	0.406	0.353	0.175	0.329

Table 7
Spearman's rank correlation results comparison to related work

Approach	IT companies	Hollywood Celebrities	Television Series	Video Games	Chuck Norris	All 21 entities
baseline	0.559	0.639	0.529	0.451	0.458	0.541
KORE	0.759	0.715	0.599	0.760	0.498	0.698
CombIC	0.644	0.690	0.643	0.532	0.558	0.624
ER	0.727	0.643	0.633	0.519	0.477	0.630
DB2Vec Similarity	0.743	0.734	0.635	0.669	0.628	0.692
DB2Vec Relatedness	0.811	0.778	0.711	0.658	0.670	0.736

Table 8

Document similarity results - Pearson’s linear correlation coefficient (r) Spearman’s rank correlation (ρ) and their harmonic mean μ

Model	r	ρ	μ
DB2vec CBOw 200 4	0.681	0.535	0.599
DB2vec CBOw 500 4	0.677	0.530	0.594
DB2vec SG 200 4	0.639	0.520	0.573
DB2vec SG 500 4	0.641	0.516	0.572
DB2vec CBOw 200 8	0.658	0.491	0.562
DB2vec CBOw 500 8	0.683	0.512	0.586
DB2vec SG 200 8	0.708	0.556	0.623
DB2vec SG 500 8	0.686	0.527	0.596
WD2vec CBOw 200 4	0.568	0.383	0.458
WD2vec CBOw 500 4	0.593	0.386	0.467
WD2vec SG 200 4	0.606	0.385	0.471
WD2vec SG 500 4	0.613	0.343	0.440
Wiki2vec	0.662	0.513	0.578

The results for the related approaches were copied from the respective papers, except for ESA, which was copied from [48], where it is calculated via public ESA REST endpoint¹². The results show that our document similarity approach outperforms all of the related approaches for both Pearson’s linear correlation coefficient and Spearman’s rank correlation, as well as their harmonic mean.

We do not compare our approach to the machine-learning approach proposed by Huang et al. [21], because that approach is a supervised one, which is tailored towards the dataset, whereas ours (as well as the others we compare to) are unsupervised.

Table 9

Comparison of the document similarity approach to the related work

Approach	r	ρ	μ
TF-IDF	0.398	0.224	0.287
AnnOv	0.590	0.460	0.517
LSA	0.696	0.463	0.556
SSA	0.684	0.488	0.570
GED	0.630	\	\
ESA	0.656	0.510	0.574
GBSS	0.704	0.519	0.598
DB2Vec	0.708	0.556	0.623

¹²<http://vmdeb20.der1.ie:8890/esaservice>

7. Recommender Systems

As discussed in Section 2.3, the *Linked Open Data* (LOD) initiative [3] has opened new interesting possibilities to realize better recommendation approaches. Given that the items to be recommended are linked to a LOD dataset, information from LOD can be exploited to determine which items are considered to be similar to the ones that the user has consumed in the past, allowing to discover hidden information and implicit relations between objects. While LOD is rich in high quality data, it is still challenging to find an effective and efficient way of exploiting the knowledge for content-based recommendations. So far, most of the proposed approaches in the literature are supervised or semi-supervised, which means cannot work without human interaction. New challenges, and at the same time new opportunities, arise from unsupervised approaches for feature learning.

RDF graph embeddings are a promising way to approach those challenges and build content-based recommender systems. As for entity similarity in the previous section, the cosine similarity between the latent vectors representing the items can be interpreted as a measure of reciprocal proximity and then exploited to produce recommendations.

In this paper, we explore the development of a hybrid RS, leveraging the latent features extracted with the proposed unsupervised learning technique. The hybrid system takes advantage of both RDF graph embeddings and Factorization Machines (FMs) [55], an effective method combining Support Vector Machines with factorization models.

7.1. Factorization Machines

Factorization Machines (FMs) [55] are a general predictor model which rely on the advantages of Support Vector Machines (SVMs) and factorization models. SVMs have been successfully applied in many classification and regression problems but they have been proved to be not effective in those settings like collaborative filtering which are characterized by huge sparsity. This happens since in sparse settings there is not enough data to estimate a “dense” factorization as SVMs would do. FM use a factorized parametrization, breaking the independence of the interaction parameters by factorizing them. The difference clearly emerges when the model equations of degree 2 for a factorization machine and for a SVM are directly compared. Being aware that in a prediction task the general

goal is to estimate a function $y : \mathcal{R}^n \rightarrow T$, from a real valued feature vector $x \in \mathcal{R}^n$ to a target domain T , the model equation [55] for a FM of degree 2 is given by:

$$\hat{y}(x) = w_0 + \sum_{i=1}^n w_i x_i + \sum_{i=1}^n \sum_{j=i+1}^n (v_i^T v_j) x_i x_j \quad (9)$$

where \hat{y} represents the estimate of function y . In (9) the parameters to be estimated are the global bias $w_0 \in \mathcal{R}$, the vector $w = [w_1, \dots, w_n] \in \mathcal{R}^n$ whose component w_i models the strength of i_{th} variable and a matrix $V \in \mathcal{R}^{n \times k}$, which describes the i_{th} variable with its i_{th} row v_i , i.e., with a vector of size k , where k is a hyperparameter defining the factorization dimensionality. On the other hand the model equation for a SVM of the same degree 2 is defined as:

$$\hat{y}(x) = w_0 + \sqrt{2} \sum_{i=1}^n w_i x_i + \sum_{i=1}^n w_{i,i}^{(2)} x_i^2 + \sqrt{2} \sum_{i=1}^n \sum_{j=i+1}^n w_{i,j}^{(2)} x_i x_j \quad (10)$$

with model parameters $w_0 \in \mathcal{R}$, $w \in \mathcal{R}^n$ and the symmetric matrix $W^{(2)} \in \mathcal{R}^{n \times n}$. From the comparison of equations (9) and (10) we can realize that both model all nested interactions up to degree 2 but with a fundamental difference: the estimation of the interactions between variables in the SVM model is done directly and independently with $w_{i,j}$ whereas it is parametrized in the FM model. In other words, parameters $w_{i,j}$ and $w_{i,l}$ will be completely independent in the SVM case while parameters $v_i^T v_j$ and $v_i^T v_l$ overlap thanks to shared parameter v_i .

The strength of FMs models is also due to linear complexity [55] and due to the fact that they can be optimized in the primal and they do not rely on support vectors like SVMs. Finally [55] shows that many of the most successful approaches for the task of collaborative filtering are de facto subsumed by FMs. It is the case for example of Matrix Factorization [71,66] and SVD++ [26].

7.2. Experiments

We evaluate different variants of our approach on three datasets, and compare them to common approaches for creating content-based item representa-

tions from LOD, as well as to state of the art collaborative and hybrid approaches. Furthermore, we investigate the use of two different LOD datasets as background knowledge, i.e., DBpedia and Wikidata.

7.2.1. Datasets

In order to test the effectiveness for the recommendation task of the graph embeddings technique previously introduced, we performed an extensive evaluation in terms of ranking accuracy on three datasets belonging to different domains, i.e., `Movielens`¹³ for movies, `LibraryThing`¹⁴ for books and `Last.fm`¹⁵ for music. The first dataset, `Movielens 1M`, contains 1 million 1-5 stars ratings from 6,040 users on 3,952 movies. The dataset `LibraryThing` contains more than 2 millions ratings from 7,279 users on 37,232 books. As in the dataset there are many duplicated ratings, when a user has rated more than once the same item, her last rating is selected. The unique ratings are 749,401, in the range from 1 to 10. Both `Movielens` and `LibraryThing` datasets contain explicit ratings and to test the approach also on implicit feedbacks, a third dataset built on the top of the `Last.fm` music system is considered. `Last.fm` contains 92,834 interactions between 1,892 users and 17,632 musical artists. Each iteration is annotated with the corresponding listening count.

The original datasets are enriched with side information thanks to the item mapping and linking to DBpedia technique detailed in [45], whose dump is available at <https://github.com/sisinflab/LODrecsys-datasets>. Since not all the items have a corresponding resource in DBpedia, after the mapping, the versions of `Movielens`, `LibraryThing` and `Last.fm` datasets contain respectively 3,883 movies, 11,695 books and 11,180 musical artists.

The datasets are finally preprocessed to guarantee a fair comparison with the state of the art approaches described in [8]. Here, the authors propose to (i) remove popularity biases from the evaluation not considering the top 1% most popular items, (ii) reduce the sparsity of `Movielens` dataset in order to have at least a sparser test dataset and (iii) remove from `LibraryThing` and `Last.fm` users with less than five ratings and items rated less than five times. The final statistics on the three datasets are reported in Table 10.

¹³<http://grouplens.org/datasets/movielens/>

¹⁴<https://www.librarything.com/>

¹⁵<http://www.lastfm.com>

Table 10
Statistics about the three datasets

	MovieLens	LibraryThing	Last.fm
Number of users	4,186	7,149	1,875
Number of items	3,196	4,541	2,432
Number of ratings	822,597	352,123	44,981
Data sparsity	93.85%	98.90%	99.01%

7.2.2. Evaluation Protocol

The ranking setting for the recommendation task consists of producing a ranked list of items to suggest to the user and in practical situations turns into the so-called top- N recommendation task, where just a cut-off of the ranked list of size N is provided to the user. This setting has recently replaced the rating prediction, because of the increasing awareness that the user is not interested in an accurate prediction of the item rating, but is looking for a (limited) list of items extracted from the pool of available ones.

As evaluation ranking protocol for our comparison, we adopted the *all unrated items* methodology presented in [72] and already used in [8]. Such methodology asks to predict a score for each item not rated by a user, irrespective of the existence of an actual rating, and to compare the recommendation list with the test set.

The metrics involved in the experimental comparison are three well-known ranking measures for recommendation accuracy, i.e., precision, recall and nDCG.

- *precision@N* [57] represents the fraction of relevant items in the top- N recommendations.
- *recall@N* [57] indicates the fraction of relevant items, in the user test set, occurring in the top- N list. As relevance threshold, we set 4 for `MovieLens` and 8 for `LibraryThing`, as previously done in [8].

Although precision and recall are good indicators to evaluate the accuracy of a recommendation engine, they are not rank-sensitive.

- The normalized Discounted Cumulative Gain *nDCG@N* [2] instead takes into account also the position in the recommendation list, being defined as

$$\text{nDCG@N} = \frac{1}{\text{iDCG}} \cdot \sum_{i=1}^N \frac{2^{\text{rel}(u,i)} - 1}{\log_2(1+i)} \quad (11)$$

where $\text{rel}(u, i)$ is a boolean function representing the relevance of item i for user u and iDCG is a normalization factor that sets nDCG@N value to 1 when an ideal ranking is returned [2].

As suggested in [72] and set up in [8], in the computation of nDCG@N we fixed a default “neutral” value for those items with no ratings, i.e., 3 for `MovieLens` and 5 for `LibraryThing`.

All the results have been computed @10, that is considering the top-10 list recommended to each user and then averaging across all users.

7.2.3. Experimental Setup

The target of this experimental section is two-fold. On the one hand, we show that the proposed graph embeddings technique outperforms other strategies for feature creation. On the other hand, we combine it with a hybrid RS and compare the resulting system’s relative performances with state of the art approaches. The first goal is pursued by implementing an item-based K-Nearest Neighbor method, hereafter denoted as `ItemKNN`, with cosine similarity among features vectors. As data mining features, the approaches based on direct relations and graph substructures, as detailed in Section 5.1, are considered.

For what concerns the comparison against some of the most promising collaborative and hybrid approaches currently available, we consider the results recently published in [8]. There, the authors show that the best effectiveness in terms of ranking accuracy is reached by the following algorithms:

- `SLIM` [42] is a Sparse Linear Method for top- N recommendation that learns a sparse coefficient matrix for the items involved in the system by only relying on the users purchase/ratings profile and by solving a L1-norm and L2-norm regularized optimization problem.
- `BPR-SSLIM` is a Sparse Linear Method using item Side information (SSLIM) [43] and the Bayesian Personalized Ranking (BPR) optimization criterion [56].
- `SPRank` [8] is a novel hybrid recommender system that solves the top- N recommendation problem in a learning to rank fashion, exploiting the freely available knowledge in the Linked Open Data to build semantic path-based features.

In [8], it has been proved that `SPRank` is able to outperform `SLIM` and `BPR-SSLIM` (at least according to Precision and nDCG) in those context characterized by higher sparsity, where reasonably the con-

tribution of the content is more essential. For the sake of completeness, it is necessary to remind that in [8] the aforementioned approaches are tested against well established collaborative algorithms for rating prediction, e.g., *Biased Matrix Factorization* [27], and for items ranking, such as *Bayesian Personalized Ranking* [56] and *Soft Margin Ranking Matrix Factorization* [78]. A pair of algorithms that work on heterogeneous information networks are also involved in the comparison. They are *PathRank* [30], an extension of Personalized PageRank for heterogeneous networks, and *HeteRec* [81] that represents the connectivity between users and items through a matrix factorization algorithm that uses meta-path based latent features.

7.3. Results

We present the results by a purely content-based RS and a hybrid RS using RDF2Vec.

7.3.1. Content-based Approach

The first experiment aims to point out the validity of the proposed graph embeddings technique for features generation in the context of content-based RS, and relies on a relatively simple recommendation algorithm, i.e., the item-based K-Nearest Neighbor approach [57] with cosine similarity. Formally, this method evaluates the closeness of items through cosine similarity between the corresponding features vectors and then selects a subset of them – the neighbors – for each item, that will be used to estimate the rating of user u for a new item i as follows:

$$r^*(u, i) = \sum_{j \in \text{ratedItems}(u)} \text{cosineSim}(j, i) \cdot r_{u,j}$$

where $\text{ratedItems}(u)$ is the set of items already evaluated by user u , $r_{u,j}$ indicates the rating for item j by user u and $\text{cosineSim}(j, i)$ is the cosine similarity score between items j and i . In our experiments, the size of the considered neighbourhood is limited to 5.

The vector representation of items made of latent features, output of the RDF graph embeddings with graph walks, is evaluated against all the other data mining features listed in Section 5.1. Tables 11, 12 and 13 contain the values of precision, recall and nDCG, respectively for *Movielens*, *LibraryThing* and *Last.fm*. The computation of recommendations

has been done with the publicly available library RankSys.¹⁶

Table 11

Results of the ItemKNN approach on *Movielens* dataset with reference to different computation of features.

Strategy	Precision	Recall	nDCG
DB2vec CBOW 200 4	0.03893	0.02167	0.30782
DB2vec CBOW 500 4	0.03663	0.02088	0.30557
DB2vec SG 200 4	0.05681	0.03119	0.31828
DB2vec SG 500 4	0.05786	0.0304	0.31726
DB2vec CBOW 200 8	0.01064	0.00548	0.29245
DB2vec CBOW 500 8	0.01137	0.00567	0.29289
DB2vec SG 200 8	0.04424	0.02693	0.30997
DB2vec SG 500 8	0.02191	0.01478	0.29863
WD2vec CBOW 200 4	0.01217	0.00596	0.29362
WD2vec CBOW 500 4	0.01027	0.00427	0.29211
WD2vec SG 200 4	0.02902	0.01479	0.30189
WD2vec SG 500 4	0.02644	0.01246	0.29967
types	0.00313	0.00145	0.28864
categories	0.0305	0.02093	0.30444
rel in	0.01122	0.00589	0.29183
rel out	0.02844	0.01607	0.30274
rel in & out	0.02852	0.01566	0.3006
rel-vals in	0.03883	0.02293	0.29411
rel-vals out	0.01279	0.00971	0.29378
rel-vals in & out	0.01174	0.00913	0.29333
WC2	0.00684	0.00343	0.29032
WL12	0.00601	0.00288	0.28977

The first conclusion that can be drawn from Tables 11, 12 and 13 is that the best approach for all datasets is retrieved with a skip-gram model and with a size of 200 for vectors built upon DBpedia. For the sake of true, on *Movielens* the highest value of precision is gained using vector size of 500, but the size 200 is prevalent according to the F1 measure, i.e., the harmonic mean of precision and recall. A substantial difference concerns the exploratory depth of the random walks, since for *Movielens* the results re-

¹⁶<http://ranksys.org/>

Table 12

Results of the ItemKNN approach on `LibraryThing` dataset with reference to different computation of features.

Strategy	Precision	Recall	nDCG
DB2vec CBOW 200 4	0.05127	0.11777	0.21244
DB2vec CBOW 500 4	0.05065	0.11557	0.21039
DB2vec SG 200 4	0.05719	0.12763	0.2205
DB2vec SG 500 4	0.05811	0.12864	0.22116
DB2vec CBOW 200 8	0.00836	0.02334	0.14147
DB2vec CBOW 500 8	0.00813	0.02335	0.14257
DB2vec SG 200 8	0.07681	0.17769	0.25234
DB2vec SG 500 8	0.07446	0.1743	0.24809
WD2vec CBOW 200 4	0.00537	0.01084	0.13524
WD2vec CBOW 500 4	0.00444	0.00984	0.13428
WD2vec SG 200 4	0.06416	0.14565	0.23309
WD2vec SG 500 4	0.06031	0.14194	0.22752
types	0.01854	0.04535	0.16064
categories	0.06662	0.15258	0.23733
rel in	0.04577	0.10219	0.20196
rel out	0.04118	0.09055	0.19449
rel in & out	0.04531	0.10165	0.20115
rel-vals in	0.06176	0.14101	0.22574
rel-vals out	0.06163	0.13763	0.22826
rel-vals in & out	0.06087	0.13662	0.22615
WC2	0.00159	0.00306	0.12858
WL12	0.00155	0.00389	0.12937

lated to a depth of 4 overtake those computed with a depth of 8, while the tendency is reversed for both `LibraryThing` and `Last.fm`. Secondly, the advantage of the Skip-Gram model over CBOW is a constant both on `DBpedia` and `Wikidata` and is particularly evident when the model involves longer random walks, i.e., with depth 8. Comparing the LOD datasets, it clearly emerges that `DBpedia` lets to gain higher values than `Wikidata` for each metric involved, but it turns out that `Wikidata` is quite effective on `LibraryThing`, where the skip-gram vectors with depth of 4 exceed the corresponding `DBpedia` vectors. Moving to the features extracted from direct relations, the contribution of the “categories” stands clearly out, together with relations-values “rel-vals”, especially when just incoming relations are considered. The use of methods based on kernels for features extraction, i.e., WC2 and WL12 approaches, seems not to provide significant advantages to the recommendation algorithm.

Table 13

Results of the ItemKNN approach on `Last.fm` with reference to different computation of features.

Strategy	Precision	Recall	nDCG
DB2vec CBOW 200 4	0.01749	0.09915	0.06435
DB2vec CBOW 500 4	0.01769	0.10016	0.06404
DB2vec SG 200 4	0.02015	0.11109	0.07232
DB2vec SG 500 4	0.02001	0.10978	0.07448
DB2vec CBOW 200 8	0.00944	0.05349	0.03311
DB2vec CBOW 500 8	0.00964	0.0563	0.03166
DB2vec SG 200 8	0.0234	0.1359	0.08719
DB2vec SG 500 8	0.02088	0.12248	0.07789
WD2vec CBOW 200 4	0.00133	0.00785	0.00382
WD2vec CBOW 500 4	0.001	0.00532	0.00408
WD2vec SG 200 4	0.00612	0.03388	0.02157
WD2vec SG 500 4	0.00658	0.03932	0.02382
types	0.00525	0.03256	0.01826
categories	0.01762	0.09889	0.06023
rel in	0.00625	0.03625	0.02042
rel out	0.00519	0.02757	0.01733
rel in & out	0.00718	0.04205	0.02567
rel-vals in	0.0185	0.10502	0.06733
rel-vals out	0.00805	0.04585	0.0248
rel-vals in & out	0.00339	0.01774	0.00982
WC2	0.00086	0.00401	0.00241
WL12	0.00086	0.00344	0.00215

To point out that the latent features built upon RDF graph are able to capture its structure, placing closely semantically similar items, some examples of the neighbouring sets retrieved using the graph embeddings technique are provided. These sets are directly exploited by the ItemKNN algorithm to produce recommendations. Table 14 is related to movies and to the strategy “DB2vec SG 200 4”, and displays that neighboring items are highly relevant and close to the query item, i.e., the item for which neighbors are searched for. Figure 4 depicts the 2D PCA projection of the movies in that table, showing that similar films are actually projected closely to each other.

7.3.2. Hybrid Approach

To second goal of our experiments is to show that a hybrid recommender system, leveraging the item content built with the graph embeddings technique, can compete against some of the most promising state of

Table 14

Examples of K-Nearest Neighbor sets on *Movielens*, for the Skip-Gram model with depth of 4 and size vectors 200, on DBpedia.

Query Movie	K Nearest Neighbours
Batman	Batman Forever, Batman Returns, Batman & Robin, Superman IV: The Quest for Peace, Dick Tracy
Bambi	Cinderella, Dumbo, 101 Dalmatians, Pinocchio, Lady and the Tramp
Star Trek: Generations	Star Trek VI: The Undiscovered Country, Star Trek: Insurrection, Star Trek III: The Search for Spock, Star Trek V: The Final Frontier, Star Trek: First Contact




Fig. 4. Two-dimensional PCA projection of the 200-dimensional Skip-Gram vectors of movies in Table 14.

the art approaches for recommendation. A hybrid RS based on Factorization Machines is utilized, namely the *ranking factorization machine* algorithm available in the *Graphlab* recommender toolkit¹⁷. This algorithm allows to define both users and items side data. Grounding on the capability to automatically cluster semantically similar items pointed out by the experiments of the previous section, the additional information added to each item, as side data, is the list of its 5 nearest neighbors. The tuning of the hyper-parameters (specifically the numbers of latent factors

¹⁷<https://dato.com/products/create/docs/graphlab.toolkits.recommender.html>

and the maximum number of iterations) has been carried out through cross-validation on a validation set obtained selecting the 15% of the ratings of each user from the training set. The same choice for cross validation has been adopted for the competing recommendations algorithms, as already done in [8].

Tables 15, 16, and 17 report, respectively for *Movielens*, *LibraryThing* and *Last.fm*, the most promising approach according to the analysis conducted with the ItemKNN algorithm in Section 7.3.1 (Tables from 11 to 13) and the competing approaches *SPRank*, *SLIM* and *BPR-SSLIM*¹⁸. For *BPR-SSLIM*, the resources directly connected to the items are used as item side data, as in [8]. The prefix “ \mathcal{H} ” is used as a notation for the hybrid RS grounding on the relative feature vectors. *SLIM* and *BPR-SSLIM* are implemented through the publicly available software library *MyMediaLite*¹⁹ [14], while for *SPRank*, the implementation provided by its authors, available at <https://github.com/sisinflab/lodreclib>, is used. We repeated the experiments three times for each datasets and then averaged the results across the three runs.

Looking at Tables 15, 16 and 17, we can observe performance improvements on all datasets and according to each of the metrics involved, with the only exception of the precision on the *Last.fm* dataset. In particular, on *Movielens*, the hybrid *SPRank* cannot compete with the collaborative method *SLIM* and to its extended version with side data, *BPR-SSLIM*, probably because of the key contribution of collaborative information on this dataset. In fact, *Movielens* is the denser dataset in the comparison, containing only experienced users with more than fifty ratings. Nevertheless, the approach “ \mathcal{H} -DB2vec SG 200 4” outperforms the competing ones, reasonably with the key contribution of factorization models in FM. On the other datasets, whose sparsity better reproduces real-world scenarios, the contribution of content appears more determinant. This emerges in particular on *LibraryThing*, where *BPR-SSLIM* and *SPRank* overtake *SLIM* and “ \mathcal{H} -DB2vec SG 200 8” improves the results even further. On the third dataset, i.e., *Last.fm*, the hybrid approaches generally perform better than *SLIM* (except on precision for

¹⁸The complete version of these tables and the code for implementing the hybrid system are available at <http://sisinflab.poliba.it/recommender-systems/hybridSystemRDF2Vec.html>

¹⁹<http://www.mymedialite.net>

BPR-SSLIM) and the approach based on graph embeddings reports a further enhancement in recall and nDCG. The results are statistically significant according to the Wilcoxon signed rank with $p < 0.001$, for less than a different specification. In conclusion, the choice of using Factorization Machines and of building item side data by means of the RDF graph embeddings technique, turns out to be favorable for fostering accuracy in general and getting improvements, on every dataset, on the rank-sensitive metric nDCG. We remember that the importance of providing to the user an opportunely *ranked* list of recommendations is nowadays recognized as fundamental in the RS research community.

Table 15

Comparative results on `MovieLens` dataset. The differences between “ \mathcal{H} -DB2vec SG 200 4” and the other methods are statistically significant according to the Wilcoxon signed rank with $p < 0.001$ (*) or $p < 0.05$ (**), with the only exception of SLIM in terms of recall.

Strategy	Precision	Recall	nDCG
\mathcal{H} -DB2vec SG 200 4	0.2698	0.1480	0.4782
SPRank	0.1613*	0.0785*	0.4056*
SLIM	0.2631*	0.1469	0.4597*
BPR-SSLIM	0.2636*	0.1439**	0.4674*

Table 16

Comparative results on `LibraryThing` dataset. The differences between “ \mathcal{H} -DB2vec SG 200 8” and the other methods are statistically significant according to the Wilcoxon signed rank with $p < 0.001$.

Strategy	Precision	Recall	nDCG
\mathcal{H} -DB2vec SG 200 8	0.1061	0.241	0.3127
SPRank	0.1019	0.2232	0.3068
SLIM	0.0543	0.0988	0.2317
BPR-SSLIM	0.0962	0.2328	0.3051

8. Conclusion and Outlook

In this paper, we have presented RDF2Vec, an approach for learning latent numerical representations of entities in RDF graphs. In this approach, we first convert the RDF graphs in a set of sequences using

Table 17

Comparative results on `Last.fm` dataset. The differences between “ \mathcal{H} -DB2vec SG 200 8” and the other methods are statistically significant according to the Wilcoxon signed rank with $p < 0.001$.

Strategy	Precision	Recall	nDCG
\mathcal{H} -DB2vec SG 200 8	0.0607	0.3408	0.2266
SPRank	0.0841	0.2043	0.1918
SLIM	0.0723	0.1548	0.1395
BPR-SSLIM	0.0465	0.2649	0.1759

two strategies, Weisfeiler-Lehman Subtree RDF Graph Kernels and graph walks, which are then used to build neural language models. The evaluation shows that such entity representations could be used in three different tasks. In each of those tasks, they were capable of outperforming standard feature generation approaches, i.e., approaches that turn (subsets of) RDF graphs into propositional models.

We have explored different variants of building embedding models. While there is no universally best performing approach, we can observe some trends. With respect to the first step of the transformation, i.e., the construction of sequences, kernel transformations lead to better results than (random) walks. However, they do not scale well to large-scale knowledge graphs, such as DBpedia or Wikidata. With respect to the second step, i.e., the actual computation of the embeddings, we have observed that Skip-Gram (SG) models in most cases outperform Continuous-Bag-of-Words (CBOW) models. The other characteristics of the models (e.g., the dimensionality of the embedding space) show less clear trends towards an optimal setting.

While constructing a vector space embedding for a large-scale knowledge graph, such as DBpedia or Wikidata, can be costly, we have shown that this step has to be taken only once, as the embeddings can be reused on various tasks. This is particularly interesting for such cross-domain knowledge graphs, which can be used in a variety of scenarios and applications.

For the moment, we have defined some constraints for the construction of the embeddings. We do not use literal values, and we do not particularly distinguish between the schema and the data level of a graph. The former constraint has some limitations, e.g., when it comes to the tasks of determining entity similarity: for example, the similarity of two movies in terms of release date and budget or the similarity of two cities in terms of area and population is currently not captured by the models. Schema level and data level similarity are currently implicitly interwoven, but in particular for knowledge graphs with richer schemas (e.g.,

YAGO with its type hierarchy of several hundred thousand types), distinguishing embeddings of the schema and data level might become beneficial.

Apart from using vector space embeddings when exploiting LOD data sources, they may also become an interesting technique for improving those sources (such as public knowledge graphs) as such. Among others, the proposed approach could also be used for link prediction, entity typing, or error detection in knowledge graphs [51], as shown in [36,41]. Similarly to the entity and document modeling, the approach can be extended for entity summarization, which is also an important task when consuming and visualizing large quantities of data [5].

Summarizing, we have shown that it is possible to adapt the technique of word embeddings to RDF graphs, and that those embeddings lead to compact vector representations of entities. We have shown that those vector representations help building approaches which outperform many state of the art tools on various tasks, e.g., data mining with background knowledge or recommender systems. Furthermore, the proposed vector space representations are *universal* in the sense that they are not task specific, i.e., a vector space embedding for a general graph like DBpedia or Wikidata can be built once and reused for several tasks. As the embeddings used in this paper are publicly available, they are a versatile asset which can be exploited for various knowledge-intensive tasks in future research.

Acknowledgements The work presented in this paper has been partly funded by the German Research Foundation (DFG) under grant number PA 2373/1-1 (Mine@LOD), and partially funded by the project PON03 PE_00136_1 Digital Services Ecosystem: DSE. Jessica Rosati also acknowledges support of I.B.M. Ph.D. fellowship 2015-2016.

References

- [1] Aggarwal, N., Buitelaar, P.: Wikipedia-based distributional semantics for entity relatedness. In: 2014 AAAI Fall Symposium Series (2014)
- [2] Bellogín, A., Cantador, I., Castells, P.: A comparative study of heterogeneous item recommendations in social systems. *Inf. Sci.* 221, 142–169 (Feb 2013), <http://dx.doi.org/10.1016/j.ins.2012.09.039>
- [3] Bizer, C., Heath, T., Berners-Lee, T.: Linked Data – The Story So Far. *International journal on semantic web and information systems* 5(3), 1–22 (2009)
- [4] Burke, R.: Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction* 12(4), 331–370 (Nov 2002), <http://dx.doi.org/10.1023/A:1021240730564>
- [5] Cheng, G., Tran, T., Qu, Y.: Relin: relatedness and informativeness-based centrality for entity summarization. In: *The Semantic Web–ISWC*, pp. 114–129 (2011)
- [6] Cheng, W., Kasneci, G., Graepel, T., Stern, D., Herbrich, R.: Automated feature generation from structured knowledge. In: *CIKM* (2011)
- [7] Deerwester, S., Dumais, S.T., Furnas, G.W., Landauer, T.K., Harshman, R.: Indexing by latent semantic analysis. *JOURNAL OF THE AMERICAN SOCIETY FOR INFORMATION SCIENCE* 41(6), 391–407 (1990)
- [8] Di Noia, T., Ostuni, V.C., Tomeo, P., Di Sciascio, E.: Sprank: Semantic path-based ranking for top-n recommendations using linked open data. *ACM Transactions on Intelligent Systems and Technology (TIST)* (2016)
- [9] Di Noia, T., Mirizzi, R., Ostuni, V.C., Romito, D.: Exploiting the web of data in model-based recommender systems. In: *Proceedings of the Sixth ACM Conference on Recommender Systems*. pp. 253–256. *RecSys '12*, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2365952.2366007>
- [10] Di Noia, T., Mirizzi, R., Ostuni, V.C., Romito, D., Zanker, M.: Linked open data to support content-based recommender systems. In: *Proceedings of the 8th International Conference on Semantic Systems*. pp. 1–8. *I-SEMANTICS '12*, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2362499.2362501>
- [11] Di Noia, T., Ostuni, V.C., Rosati, J., Tomeo, P., Di Sciascio, E., Mirizzi, R., Bartolini, C.: Building a relatedness graph from linked open data: A case study in the IT domain. *Expert Syst. Appl.* 44, 354–366 (2016)
- [12] Fanizzi, N., d’Amato, C.: A declarative kernel for alc concept descriptions. In: *Foundations of Intelligent Systems*, pp. 322–331 (2006)
- [13] Gabrilovich, E., Markovitch, S.: Computing semantic relatedness using wikipedia-based explicit semantic analysis. In: *IJ-cAI*. vol. 7, pp. 1606–1611 (2007)
- [14] Gantner, Z., Rendle, S., Freudenthaler, C., Schmidt-Thieme, L.: Mymedialite: A free recommender system library. In: *Proceedings of the Fifth ACM Conference on Recommender Systems*. pp. 305–308. *RecSys '11*, ACM, New York, NY, USA (2011), <http://doi.acm.org/10.1145/2043932.2043989>
- [15] de Gemmis, M., Lops, P., Musto, C., Fedelucio, N., Semeraro, G.: Semantics-aware content-based recommender systems. In: Ricci, F., Rokach, L., Shapira, B. (eds.) *Recommender Systems Handbook*, pp. 119–159. Springer, 2nd edn. (2015)
- [16] Grbovic, M., Radosavljevic, V., Djuric, N., Bhamidipati, N., Savla, J., Bhagwan, V., Sharp, D.: E-commerce in your inbox: Product recommendations at scale. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1809–1818. *KDD '15*, ACM, New York, NY, USA (2015), <http://doi.acm.org/10.1145/2783258.2788627>
- [17] Harris, Z.S.: *Mathematical Structures of Language*. Wiley, New York, NY, USA (1968)
- [18] Hassan, S., Mihalcea, R.: Semantic relatedness using salient semantic analysis. In: *AAAI* (2011)
- [19] Hassanzadeh, O., Consens, M.: M.: Linked movie data base. In: *Workshop on Linked Data on the Web* (2009)

- [20] Hoffart, J., Seufert, S., Nguyen, D.B., Theobald, M., Weikum, G.: Kore: keyphrase overlap relatedness for entity disambiguation. In: Proceedings of the 21st ACM international conference on Information and knowledge management. pp. 545–554. ACM (2012)
- [21] Huang, L., Milne, D., Frank, E., Witten, I.H.: Learning a concept-based document similarity measure. *Journal of the American Society for Information Science and Technology* 63(8), 1593–1608 (2012)
- [22] Huang, Y., Tresp, V., Nickel, M., Kriegel, H.P.: A scalable approach for statistical learning in semantic graphs. *Semantic Web* (2014)
- [23] Hulpuş, I., Prangnawarat, N., Hayes, C.: Path-based semantic relatedness on linked data and its use to word and entity disambiguation. In: International Semantic Web Conference. pp. 442–457. Springer (2015)
- [24] Kappara, V.N.P., Ichise, R., Vyas, O.: Liddm: A data mining system for linked data. In: LDOW (2011)
- [25] Khan, M.A., Grimnes, G.A., Dengel, A.: Two pre-processing operators for improved learning from semanticweb data. In: RCOMM (2010)
- [26] Koren, Y.: Factorization meets the neighborhood: A multifaceted collaborative filtering model. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 426–434. KDD '08, ACM, New York, NY, USA (2008), <http://doi.acm.org/10.1145/1401890.1401944>
- [27] Koren, Y., Bell, R., Volinsky, C.: Matrix factorization techniques for recommender systems. *Computer* 42(8), 30–37 (Aug 2009), <http://dx.doi.org/10.1109/MC.2009.263>
- [28] Kramer, S., Lavrač, N., Flach, P.: Propositionalization approaches to relational data mining. In: Relational Data Mining, pp. 262–291. Springer Berlin Heidelberg (2001)
- [29] Lee, M., Pincombe, B., Welsh, M.: An empirical evaluation of models of text document similarity. *Cognitive Science Society* (2005)
- [30] Lee, S., Park, S., Kahng, M., Lee, S.g.: Pathrank: A novel node ranking measure on a heterogeneous graph for recommender systems. In: Proceedings of the 21st ACM International Conference on Information and Knowledge Management. pp. 1637–1641. CIKM '12, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2396761.2398488>
- [31] Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morse, M., van Kleef, P., Auer, S., Bizer, C.: DBpedia – A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web Journal* (2013)
- [32] Lösch, U., Bloehdorn, S., Rettinger, A.: Graph kernels for rdf data. In: The Semantic Web: Research and Applications, pp. 134–148. Springer (2012)
- [33] Middleton, S.E., Roue, D.D., Shadbolt, N.R.: Ontology-Based Recommender Systems, pp. 779–796. Springer Berlin Heidelberg, Berlin, Heidelberg (2009), http://dx.doi.org/10.1007/978-3-540-92673-3_35
- [34] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013)
- [35] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Advances in neural information processing systems. pp. 3111–3119 (2013)
- [36] Minervini, P., Fanizzi, N., d’Amato, C., Esposito, F.: Scalable learning of entity and predicate embeddings for knowledge graph completion. In: 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA). pp. 162–167. IEEE (2015)
- [37] Musto, C., Semeraro, G., De Gemmis, M., Lops, P.: Word embedding techniques for content-based recommender systems: an empirical evaluation. In: RecSys Posters, ser. CEUR Workshop Proceedings, P. Castells, Ed. vol. 1441
- [38] Musto, C., Semeraro, G., Lops, P., de Gemmis, M.: Random Indexing and Negative User Preferences for Enhancing Content-Based Recommender Systems, pp. 270–281. Springer Berlin Heidelberg, Berlin, Heidelberg (2011), http://dx.doi.org/10.1007/978-3-642-23014-1_23
- [39] Musto, C., Semeraro, G., Lops, P., de Gemmis, M.: Contextual eVSM: A Content-Based Context-Aware Recommendation Framework Based on Distributional Semantics, pp. 125–136. Springer Berlin Heidelberg, Berlin, Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-39878-0_12
- [40] Mynarz, J., Svátek, V.: Towards a benchmark for LOD-enhanced knowledge discovery from structured data. In: The Second International Workshop on Knowledge Discovery and Data Mining Meets Linked Open Data (2013)
- [41] Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E.: A review of relational machine learning for knowledge graphs: From multi-relational link prediction to automated knowledge graph construction. *arXiv preprint arXiv:1503.00759* (2015)
- [42] Ning, X., Karypis, G.: SLIM: sparse linear methods for top-n recommender systems. In: 11th IEEE International Conference on Data Mining, ICDM 2011, Vancouver, BC, Canada, December 11-14, 2011. pp. 497–506 (2011), <http://dx.doi.org/10.1109/ICDM.2011.134>
- [43] Ning, X., Karypis, G.: Sparse linear methods with side information for top-n recommendations. In: Proceedings of the Sixth ACM Conference on Recommender Systems. pp. 155–162. RecSys '12, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2365952.2365983>
- [44] Nunes, B.P., Kawase, R., Fetahu, B., Dietze, S., Casanova, M.A., Maynard, D.: Interlinking documents based on semantic graphs. *Procedia Computer Science* 22, 231–240 (2013)
- [45] Ostuni, V.C., Noia, T.D., Sciascio, E.D., Mirizzi, R.: Top-n recommendations from implicit feedback leveraging linked open data. In: ACM RecSys '13. pp. 85–92 (2013)
- [46] Ostuni, V.C., Oramas, S., Di Noia, T., Serra, X., Di Sciascio, E.: Sound and music recommendation with knowledge graphs. *ACM Transactions on Intelligent Systems and Technology (TIST)* (2016), to appear
- [47] Ozsoy, M.G.: From word embeddings to item recommendation. *arXiv preprint arXiv:1601.01356* (2016)
- [48] Paul, C., Rettinger, A., Mogadala, A., Knoblock, C.A., Szekely, P.: Efficient graph-based document similarity. In: International Semantic Web Conference. pp. 334–349. Springer (2016)
- [49] Paulheim, H.: Exploiting linked open data as background knowledge in data mining. In: Workshop on Data Mining on Linked Open Data (2013)
- [50] Paulheim, H.: Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web (Preprint)*, 1–20 (2016)

- [51] Paulheim, H.: Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods. *Semantic Web Journal* (2016), to appear
- [52] Paulheim, H., Fümkrantz, J.: Unsupervised generation of data mining features from linked open data. In: *Proceedings of the 2nd international conference on web intelligence, mining and semantics*. p. 31. ACM (2012)
- [53] Paulheim, H., Ristoski, P., Mitichkin, E., Bizer, C.: Data mining with background knowledge from the web. *RapidMiner World* (2014)
- [54] Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 701–710. ACM (2014)
- [55] Rendle, S.: Factorization machines with libfm. *ACM Trans. Intell. Syst. Technol.* 3(3), 57:1–57:22 (May 2012), <http://doi.acm.org/10.1145/2168752.2168771>
- [56] Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: Bpr: Bayesian personalized ranking from implicit feedback. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. pp. 452–461. UAI '09, Arlington, Virginia, United States (2009), <http://dl.acm.org/citation.cfm?id=1795114.1795167>
- [57] Ricci, F., Rokach, L., Shapira, B., Kantor, P.B.: *Recommender Systems Handbook*. Springer-Verlag New York, Inc., New York, NY, USA, 2nd edn. (2015)
- [58] Ristoski, P., Bizer, C., Paulheim, H.: Mining the web of linked data with rapidminer. *Web Semantics: Science, Services and Agents on the World Wide Web* 35, 142–151 (2015)
- [59] Ristoski, P., Mencía, E.L., Paulheim, H.: A hybrid multi-strategy recommender system using linked open data. In: *Semantic Web Evaluation Challenge*. pp. 150–156. Springer (2014)
- [60] Ristoski, P., Paulheim, H.: A comparison of propositionalization strategies for creating features from linked open data. In: *Linked Data for Knowledge Discovery* (2014)
- [61] Ristoski, P., Paulheim, H.: Rdf2vec: Rdf graph embeddings for data mining. In: *International Semantic Web Conference (To Appear)*. Springer (2016)
- [62] Ristoski, P., Paulheim, H.: Semantic web in data mining and knowledge discovery: A comprehensive survey. *Web Semantics: Science, Services and Agents on the World Wide Web* (2016)
- [63] Ristoski, P., de Vries, G.K.D., Paulheim, H.: A collection of benchmark datasets for systematic evaluations of machine learning on the semantic web. In: *International Semantic Web Conference (To Appear)*. Springer (2016)
- [64] Rosati, J., Ristoski, P., Noia, T.D., Leone, R.D., Paulheim, H.: Rdf graph embeddings for content-based recommender systems. In: *Proceedings of the 3rd Workshop on New Trends in Content-based Recommender Systems (CBRecSys 2016)* (September 2016)
- [65] Sahlgrén, M.: An introduction to random indexing. In: *In Methods and Applications of Semantic Indexing Workshop at the 7th International Conference on Terminology and Knowledge Engineering, TKE 2005* (2005)
- [66] Salakhutdinov, R., Mnih, A.: Bayesian probabilistic matrix factorization using markov chain monte carlo. In: *Proceedings of the 25th International Conference on Machine Learning*. pp. 880–887. ICML '08, ACM, New York, NY, USA (2008), <http://doi.acm.org/10.1145/1390156.1390267>
- [67] Schmachtenberg, M., Bizer, C., Paulheim, H.: Adoption of the linked data best practices in different topical domains. In: *The Semantic Web–ISWC* (2014)
- [68] Schuhmacher, M., Ponzetto, S.P.: Knowledge-based graph document modeling. In: *Proceedings of the 7th ACM international conference on Web search and data mining*. pp. 543–552. ACM (2014)
- [69] Semeraro, G., Lops, P., Basile, P., de Gemmis, M.: Knowledge infusion into content-based recommender systems. In: *Proceedings of the Third ACM Conference on Recommender Systems*. pp. 301–304. RecSys '09, ACM, New York, NY, USA (2009), <http://doi.acm.org/10.1145/1639714.1639773>
- [70] Shervashidze, N., Schweitzer, P., Van Leeuwen, E.J., Mehlhorn, K., Borgwardt, K.M.: Weisfeiler-lehman graph kernels. *The Journal of Machine Learning Research* 12, 2539–2561 (2011)
- [71] Srebro, N., Rennie, J.D.M., Jaakola, T.S.: Maximum-margin matrix factorization. In: *Advances in Neural Information Processing Systems* 17. pp. 1329–1336. MIT Press (2005)
- [72] Steck, H.: Evaluation of recommendations: Rating-prediction and ranking. In: *Proceedings of the 7th ACM Conference on Recommender Systems*. pp. 213–220. RecSys '13, ACM, New York, NY, USA (2013), <http://doi.acm.org/10.1145/2507157.2507160>
- [73] Thiagarajan, R., Manjunath, G., Stumptner, M.: Computing semantic similarity using ontologies. HP Laboratories). Technical report HPL-2008-87 (2008)
- [74] Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Communications of the ACM* 57(10), 78–85 (2014)
- [75] de Vries, G.K.D.: A fast approximation of the Weisfeiler-Lehman graph kernel for RDF data. In: *ECML/PKDD* (1) (2013)
- [76] de Vries, G.K.D., de Rooij, S.: A fast and simple graph kernel for rdf. In: *DMLOD* (2013)
- [77] de Vries, G.K.D., de Rooij, S.: Substructure counting graph kernels for machine learning from rdf data. *Web Semantics: Science, Services and Agents on the World Wide Web* 35, 71–84 (2015)
- [78] Weimer, M., Karatzoglou, A., Smola, A.: Improving maximum margin matrix factorization. *Mach. Learn.* 72(3), 263–276 (Sep 2008), <http://dx.doi.org/10.1007/s10994-008-5073-7>
- [79] Widdows, D.: Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In: *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics - Volume 1*. pp. 136–143. ACL '03, Association for Computational Linguistics, Stroudsburg, PA, USA (2003), <http://dx.doi.org/10.3115/1075096.1075114>
- [80] Yanardag, P., Vishwanathan, S.: Deep graph kernels. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1365–1374. ACM (2015)
- [81] Yu, X., Ren, X., Sun, Y., Gu, Q., Sturt, B., Khandelwal, U., Norick, B., Han, J.: Personalized entity recommendation: A heterogeneous information network approach. In: *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*. pp. 283–292. WSDM '14, ACM, New York, NY, USA (2014)