

Factorization Machines Leveraging Lightweight Linked Open Data-Enabled Features for Top-N Recommendations

Guangyuan Piao^(✉) and John G. Breslin

Insight Centre for Data Analytics, National University of Ireland Galway,
IDA Business Park, Lower Dangan, Galway, Ireland
guangyuan.piao@insight-centre.org, john.breslin@nuigalway.ie

Abstract. With the popularity of Linked Open Data (LOD) and the associated rise in freely accessible knowledge that can be accessed via LOD, exploiting LOD for recommender systems has been widely studied based on various approaches such as *graph-based* or using different machine learning models with LOD-enabled features. Many of the previous approaches require construction of an additional graph to run graph-based algorithms or to extract path-based features by combining user-item interactions (e.g., likes, dislikes) and background knowledge from LOD. In this paper, we investigate *Factorization Machines* (FMs) based on particularly *lightweight* LOD-enabled features which can be directly obtained via a public SPARQL Endpoint without any additional effort to construct a graph. Firstly, we aim to study whether using FM with these *lightweight* LOD-enabled features can provide competitive performance compared to a learning-to-rank approach leveraging LOD as well as other well-established approaches such as kNN-item and BPRMF. Secondly, we are interested in finding out to what extent each set of LOD-enabled features contributes to the recommendation performance. Experimental evaluation on a standard dataset shows that our proposed approach using FM with lightweight LOD-enabled features provides the best performance compared to other approaches in terms of five evaluation metrics. In addition, the study of the recommendation performance based on different sets of LOD-enabled features indicate that *property-object lists* and *PageRank scores* of items are useful for improving the performance, and can provide the best performance through using them together for FM. We observe that *subject-property lists* of items does not contribute to the recommendation performance but rather decreases the performance.

1 Introduction

The term Linked Data, indicates a new generation of technologies responsible for the evolution of the current Web from a Web of interlinked documents to a Web of interlinked data [8]. Thanks to the Semantic Web's growth and the more recent Linked Open Data (LOD) initiative [1], a large amount of

RDF¹ data has been published in freely accessible datasets. These datasets are connected with each other to form the so-called Linked Open Data cloud². DBpedia [16] which is a 1st-class citizen in this cloud, has become one of the most important and interlinked datasets on the LOD cloud. DBpedia provides cross-domain background knowledge about entities which can be accessible via its SPARQL Endpoint³. For example, Fig. 1 shows pieces of background knowledge about the movie `dbr:The_Godfather` in RDF triples, which can be obtained from DBpedia. A RDF triple consists of a subject, a property and an object. As we can see from the figure, there can be incoming knowledge, e.g., `dbr:Carlo_Savina`→`dbo:knownFor`→`dbr:The_Godfather` where `dbr:The_Godfather` is used as an object, as well as outgoing knowledge such as `dbr:The_Godfather`→`dbo:director`→`dbr:Francis_Ford_Coppola` where `dbr:The_Godfather` is a subject. In the context of the great amount of freely accessible information, many researches have been conducted in order to consume the knowledge provided by LOD for adaptive systems such as recommender systems [2, 6].




Fig. 1. An example of background knowledge about the movie `dbr:The_Godfather` from DBpedia.

There have been many approaches for LOD-enabled recommender systems (LODRecSys) such as *semantic similarity/distance* measures, *graph-based* approaches, and *learning-to-rank* approaches by consuming LOD-enabled features. Some previous studies compared their LODRecSys approaches against well-established *collaborative filtering* approaches such as kNN and *matrix factorization* models such as BPRMF [30], and have shown the benefits of consuming background knowledge powered by LOD. On the other hand, *matrix factorization* models such as BPRMF, which do not exploit LOD-enabled features, have shown competitive performance even compared to some LODRecSys

¹ <https://www.w3.org/RDF/>.

² <http://lod-cloud.net/>.

³ <http://dbpedia.org/sparql>.

⁴ The prefix `dbr` denotes for <http://dbpedia.org/resource/>.

⁵ The prefix `dbo` denotes for <http://dbpedia.org/ontology/>.

approaches [17,20]. This has in turn motivated us to investigate factorization models consuming LOD-enabled features.

In this paper, we investigate the use of Factorization Machines (FMs), which can mimic other well-known factorization models such as *matrix factorization*, by leveraging LOD-enabled features. Previous works require increased effort to maintain an additional graph based on user-item interactions and background knowledge about items from LOD in their approaches (We will discuss this in detail in Sect. 2). In this work, we especially focus on lightweight LOD-enabled features for FM. We define *lightweight LOD features* as features that can be directly obtained via a public SPARQL Endpoint.

The contributions of this work are summarized as follows.

- We investigate lightweight LOD-enabled features, which can be directly obtained via the public DBpedia Endpoint, for FM to provide the top- N recommendations. Therefore, there is no need to construct a graph which combines user-item interactions (e.g., likes, dislikes) and background knowledge about items. In addition, we investigate to what extent different sets of these features contribute to FM in terms of recommendation performance.
- We comprehensively evaluate our approach by comparing it to other approaches such as PopRank, kNN, BPRMF, and a state-of-the-art LODRec-Sys approach SPRank [20] in terms of five different evaluation metrics.

The organization of the rest of the paper is as follows. Section 2 gives some related work, and Sect. 3 describes our proposed approach using FM with lightweight LOD-enabled features. In Sect. 4, we describe our experimental setup including the dataset and evaluation metrics. Experimental results are presented in Sect. 5. Finally, Sect. 6 concludes the paper with some brief ideas for future work.

2 Related Work

The first attempts to leverage LOD for recommender systems were by [10,25]. Heitmann et al. [10] proposed a framework using LOD for open collaborative recommender systems. The Linked Data Semantic Distance (LSD) measure [25] was one of the first works to use LOD for recommender systems in the music domain [24]. This distance measure considers direct links between two entities/nodes. In addition, it also considers that the same incoming and outgoing nodes via the same properties of two nodes in a graph such as DBpedia. Piao et al. [26,27] extended LSD by investigating different normalization strategies for the paths between two entities. These measures have been designed to work directly on LOD without considering the collaborative view of users. Based on the nature of the graph structure of DBpedia, *graph-based* approaches have been proposed [17,19]. For instance, Musto et al. [17] presented a *personalized PageRank* algorithm [7] using LOD-enabled features for the top- N recommendations. Nguyen et al. [19] investigated *SimRank* [12] and *PageRank*, and their

performance for computing similarity between entities in RDF graphs and investigated their usage to feed a content-based recommender system. Di Noia et al. [3] adapted the Vector Space Model (VSM) to a LOD-based setting, and represented the whole RDF graph as a matrix. On top of the VSM representation, they used the Support Vector Machine (SVM) as a classifier to predict if a user would like an item or not. Using the same representation, they also proposed to assign a weight to each property that represents its worth with respect to the user profile [4]. In this regard, they used a Genetic Algorithm (GA) to learn the weights of properties that minimize the misclassification errors. More recently, Di Noia et al. [20, 22] proposed SPRank, which is a semantic path-based approach using learning-to-rank algorithms. This approach first constructed a graph based on user-item interactions and the background knowledge of items from LOD. Afterwards, features, called *semantic paths*, were extracted based on the number of paths between a user and an item with min-max normalization. The extracted features were then fed into existing learning-to-rank algorithms such as LMART [33] provided by RankLib⁶. The common requirement for *graph-based* approaches as well as SPRank is that a graph has to be built based on user-item interactions and background knowledge from LOD. Our approach is different as we only consider lightweight LOD-enabled features which can be directly obtained through a public SPARQL Endpoint, and without any additional effort to build a graph. This also makes our model consume updated background knowledge of DBpedia easier when compared to other approaches such as graph-based ones which require downloading a DBpedia dump and building a graph by adding user-item interactions.

There have also been some other interesting directions related to LOD-enabled recommender systems such as the practical LODRecSys [21], explaining using LOD [18], rating predictions based on matrix factorization with semantic categories [31], and cross-domain recommendations [9, 11]. For example, Oliveira et al. [21] presented a recommender system in the movie domain that consumes LOD (not restricted to DBpedia), which was evaluated by comparing to seevl (ISWC challenge winner at 2011). Different types of evaluation metrics have been used such as accuracy, novelty etc. The authors from [18] presented ExpLOD - a framework which can generate explanations in natural language based on LOD cloud. Musto et al. [17] investigated various feature (property) selection strategies and their influences on recommendation performance in terms of accuracy and diversity in movie and book domains. Lalithsena et al. [14] proposed a novel approach using *type-* and *path-based* methods to extract a subgraph for domain specific recommendation systems. They presented that their approach can decrease 80% of the graph size without losing accuracy in the context of recommendation systems in movie and book domains. These, although interesting, are however beyond the scope of this paper and we aim to explore them in future work.

⁶ <https://sourceforge.net/p/lemur/wiki/RankLib/>.

3 Proposed Method

In this section, we first briefly introduce FMs and the optimization criteria we used in this study (Sect. 3.1). Next, we will describe our features from user-item interactions as well as background knowledge from DBpedia (Sect. 3.2).

3.1 Factorization Machines

Factorization Machines (FMs) [28], which can mimic other well known factorization models such as *matrix factorization*, *SVD++* [13], have been widely used for collaborative filtering tasks [29]. FMs are able to incorporate the high-prediction accuracy of factorization models and flexible feature engineering. An important advantage of FMs is the model equation

$$\hat{y}^{FM}(x) = w_0 + \sum_{i=1}^p w_i x_i + \sum_{i=1}^p \sum_{j>i}^p \langle v_i, v_j \rangle x_i x_j \quad (1)$$

where $w_0 \in \mathbb{R}$, x and $w \in \mathbb{R}^p$, $v_i \in \mathbb{R}^m$. The first part of the FM model captures the interactions of each input variable x_i , while the second part of it models all pairwise interactions of input variables $x_i x_j$. Each variable x_i has a latent factor v_i , which is a m -dimensional vector allows FMs work well even in highly sparse data.

Optimization. In this work, we use a *pairwise* optimization approach - Bayesian Personalized Ranking (BPR). The loss function was proposed by Rendle et al. [30].

$$l(x_1, x_2) = \sum_{x_1 \in C_u^+} \sum_{x_2 \in C_u^-} (-\log[\delta(\hat{y}^{FM}(x_1) - \hat{y}^{FM}(x_2))]) \quad (2)$$

where δ is a sigmoid function: $\delta(x) = \frac{1}{1+e^{-x}}$, and C_u^+ and C_u^- denote the set of positive and negative feedbacks respectively. L2-regularization is used for the loss function.

Learning. We use the well-known *stochastic gradient descent* algorithm to learn the parameters in our model. To avoid overfitting on the training dataset, we adopt an early stopping strategy as follows.

1. Split the dataset into training and validation sets.
2. Measure the current loss on the validation set at the end of each epoch.
3. Stop and remember the epoch if the loss has increased.
4. Re-train the model using the whole dataset.

3.2 Features

Figure 2 presents the overview of features for our FM. The details of each set of features are described below.




Fig. 2. Overview of features for Factorization Machine. PO denotes all property-objects, and SP denotes all subject-property for items in the dataset. PR denotes the PageRank scores of items.

User and Item Index. The first two sets of features indicate the indexes of the user and item in a training example. A feature value equals 1 for the corresponding user/item index, e.g., $val(U_i) = 1$ and $val(I_j) = 1$ denote an example about the i -th user and j -th item.

Property-Object List (PO). This set of features denotes all property-objects of an item i when i is a subject in RDF triples. This set of features can be obtained easily by using a SPARQL query as shown below via the DBpedia SPARQL Endpoint.

```
PREFIX dbo:<http://dbpedia.org/ontology/>
PREFIX dct:<http://purl.org/dc/terms/>
```

```
SELECT DISTINCT ?p ?o WHERE { { <itemURI> ?p ?o .
FILTER REGEX(STR(?p), '^http://dbpedia.org/ontology') } .
FILTER (STR(?p) NOT IN (dbo:wikiPageRedirects,
dbo:wikiPageExternalLink)) . FILTER ISURI(?o) }
UNION { <itemURI> ?p ?o . FILTER ( STR(?p) IN (dct:subject) ) } }
```

An intuitive way of giving feature values for PO might be to assign 1 for all property-objects of an item i (PO_i). However, it can be biased as some entities in DBpedia have a great number of property-objects while others do not. Therefore, we normalize the feature values of PO_i based on the size of PO_i so that all the feature values of PO_i sum up to 1. Formally, the feature value of j -th property-object for an item i is measured as $val(PO_i(j)) = \frac{1}{|PO_i|}$. Take the graph in Fig. 1 as an example, as we have two property-objects for the movie `dbr:The_Godfather`, each property-object of the movie will have a feature value of 0.5, respectively (see Fig. 3).

Subject-Property List (SP). Similar to the PO, we can obtain incoming background knowledge about an item i where i is an object in RDF triples. This set of features can be obtained by using a SPARQL query as shown below.

...	0.5	...	0.5	...
...	dbo:director→dbr:Francis_Ford_Coppola	...	dc:subject→dbc:Gangster_films	...

Fig. 3. An example for PO values for the movie `dbr:The_Godfather` in Fig. 1.

```
PREFIX dbo:<http://dbpedia.org/ontology/>
```

```
SELECT DISTINCT ?s ?p WHERE { ?s ?p <itemURI> .
FILTER REGEX(STR(?p), '^http://dbpedia.org/ontology') .
FILTER (STR(?p) NOT IN (dbo:wikiPageRedirects,
dbo:wikiPageExternalLink, dbo:wikiPageDisambiguates) }
```

In the same way as we normalized feature values of PO_i for an item i , we normalize the feature values of SP_i based on the size of SP_i so that all the feature values of SP_i sum up to 1. The feature value of the j -th SP for an item i is measured as $val(SP_i(j)) = \frac{1}{|SP_i|}$.

PageRank Score (PR). PageRank [23] is a popular algorithm with the purpose of measuring the relative importance of a node in a graph. In order to capture the importance of an entity in Wikipedia/DBpedia, Thalhammer et al. [32] proposed providing PageRank scores of all DBpedia entities, which are based on links using `dbo:wikiPageWikiLink` among entities. A PageRank score of an item (entity) might be a good indicator of the importance of an entity for recommendations in our case. The PageRank score of a DBpedia entity can be obtained by using the SPARQL as shown below.

```
PREFIX rdf:<http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX dbo:<http://dbpedia.org/ontology/>
PREFIX vrank:<http://purl.org/voc/vrank#>

SELECT ?score FROM <http://dbpedia.org>
FROM <http://people.aifb.kit.edu/ath/#DBpedia_PageRank>
WHERE { <itemURI> vrank:hasRank/vrank:rankValue ?score . }
```

The scale of PageRank scores is different from other feature values, which can delay the convergence of learning parameters for our model. In this regard, we normalize the PageRank scores by their maximum value.

$$val(PR_i) = \frac{PageRank_i}{\max(PageRank_j, j \in I)} \quad (3)$$

where $PageRank_i$ denotes the original PageRank score of i which is obtained from the SPARQL Endpoint, and $\max(PageRank_j, j \in I)$ denotes the maximum PageRank score of all items.

4 Experimental Setup

In this section, we introduce the dataset for our experiment (Sect. 4.1) and five evaluation metrics for evaluating the performance of the recommendations (Sect. 4.2). Afterwards, we describe four methods that have been used for comparison with our approach for evaluation (Sect. 4.3).

4.1 Dataset

We used the Movielens dataset from [20]. The dataset was originally from the Movielens dataset⁷, which consists of users and their ratings about movie items. To facilitate LODRecSys, each of the items in this dataset has been mapped into DBpedia entities if there is a mapping available⁸. In the same way as [20], we consider ratings higher than 3 as positive feedback and others as negative one. Table 1 shows details about the dataset. The dataset consists of 3,997 users and 3,082 items with 695,842 ratings where 56% of them are positive ratings. We split the dataset into training (80%) and test (20%) sets for our experiment.

Table 1. Movielens dataset statistics

# of users	3,997
# of items	3,082
# of ratings	695,842
avg. # of ratings	174
sparsity	94.35%
% of positive ratings	56%

4.2 Evaluation Metrics

We use five different evaluation metrics to measure the quality of recommendations provided by different approaches.

- **P@N**: The Precision at rank N represents the mean probability that retrieved items within the top- N recommendations are relevant to the user.

$$P@N = \frac{|\{relevant\ items\} \cap \{\text{retrieved items}@n\}|}{|\{\text{retrieved items}\}|} \quad (4)$$

- **R@N**: The Recall at rank N represents the mean probability that relevant items are successfully retrieved within the top- N recommendations.

$$R@N = \frac{|\{relevant\ items\} \cap \{\text{retrieved items}@n\}|}{|\{relevant\ items\}|} \quad (5)$$

⁷ <https://grouplens.org/datasets/movielens/1m/>.

⁸ <http://sisinflab.poliba.it/semanticweb/lod/recsys/datasets/>.

- **nDCG@N**: Precision and recall consider the relevance of items only. On the other hand, nDCG takes into account the relevance of items as well as their rank positions.

$$nDCG@N = \frac{1}{IDCG@N} \sum_{k=1}^N \frac{2^{\hat{r}_{uk}} - 1}{\log_2(1+k)} \quad (6)$$

Here, \hat{r}_{uk} denotes the rating given by a user u to the item in position k in the top- N recommendations, and $IDCG@N$ denotes the score obtained by an ideal or perfect top- N ranking and acts as a normalization factor.

- **MRR**: The Mean Reciprocal Rank (MRR) indicates at which rank the first link *relevant* to the user occurs (denoted by $rank_k$) on average.

$$MRR = \frac{1}{|U|} \sum_{k=1}^{|U|} \frac{1}{rank_k} \quad (7)$$

- **MAP**: The Mean Average Precision (MAP) measures the average of the average precision (AP) of all liked items for all users. For each user, the average precision of the user is defined as:

$$AP = \frac{\sum_{n=1}^N P@n \times like(n)}{|I|} \quad (8)$$

where n is the number of items, $|I|$ is the number of liked items of the user, and $like(n)$ is a binary function to indicate whether the user prefers the n -th item or not.

The *bootstrapped paired t-test*, which is an alternative to the paired t-test when the assumption of normality of the method is in doubt, is used for testing the significance where the significance level was set to 0.01 unless otherwise noted.

4.3 Compared Methods

We use four approaches including a baseline PopRank and other methods which have been frequently used in the literature [17,20] to evaluate our proposed method.

- **PopRank**: This is a non-personalized baseline approach which recommends items based on the popularity of each item.
- **kNN-item**: This is a collaborative filtering approach based on the k most similar items. We use a MyMedialite [5] implementation for this baseline where $k = 80$.
- **BPRMF** [30]: This is a matrix factorization approach for learning latent factors for users and items. We use a MyMedialite [5] implementation for this baseline where the dimensionality of the factorization $m = 200$.

- **SPRank** [20]: This is a *learning-to-rank* approach for LODRecSys based on *semantic paths* extracted from a graph including user-item interactions (e.g., likes, dislikes, etc.) as well as the background knowledge obtained from DBpedia. In detail, *semantic paths* are sequences of properties including *likes* and *dislikes* based on user-item interactions. For example, given the graph information `user1→likes→item1→p1→item2`, a semantic path (`likes`, `p1`) can be extracted from `user1` to `item2`. Another difference between SPRank [20] and our approach in terms of features is that the authors considered property-objects for each item including the property `dbo:wikiPageWikiLink` which cannot be queried via the DBpedia Endpoint but requires settings up a local endpoint using a DBpedia dump. On the other hand, we only considers sets of LOD-enabled features which can be obtained from a public DBpedia Endpoint. We use LMART [33] as the learning algorithm for SPRank as this approach overall provides the best performance compared to other learning-to-rank algorithms in [20]. We used the author’s implementation⁹ which has been optimized for nDCG@10.

5 Results

In this section, we first compare our approach to the aforementioned methods in terms of five evaluation metrics (Sect. 5.1). We denote our approach as LODFM, and the results of LODFM are based on best tuned parameters, i.e., $m = 200$ using PO and PR as LOD-enabled features. We discuss self comparison by using different sets of features, as well as different dimensionality m for factorization, in detail in Sect. 5.2.

Table 2. Recommendation performance compared to baselines in terms of five different evaluation metrics. The best performing strategy is in bold.

	PopRank	kNN-item	BPRMF	SPRank	LODFM
MRR	0.4080	0.5756	0.5906	0.3013	0.6218
MAP	0.1115	0.2037	0.2018	0.0612	0.2318
nDCG@1	0.2459	0.4086	0.4269	0.1758	0.4685
P@1	0.2459	0.4086	0.4269	0.1758	0.4685
R@1	0.0064	0.0132	0.0258	0.0082	0.0268
nDCG@5	0.2809	0.4049	0.4176	0.2195	0.4537
P@5	0.2240	0.3538	0.3393	0.1287	0.3829
R@5	0.0305	0.0553	0.0977	0.0291	0.1052
nDCG@10	0.3664	0.4753	0.5000	0.2845	0.5231
P@10	0.2104	0.3179	0.2883	0.1068	0.3256
R@10	0.0580	0.0978	0.1602	0.0488	0.1730

⁹ <https://github.com/sisinflab/lodreclib>.

5.1 Comparison with Baselines

The results of comparing our proposed approach with the baselines are presented in Table 2 in terms of MRR, MAP, nDCG@ N , P@ N and R@ N .

Overall, LODFM provides the best performance in terms of all evaluation metrics. In line with the results from [20], SPRank does not perform as well on the Movielens dataset compared to other collaborative filtering approaches such as kNN and BPRMF. On the other hand, we observe that LODFM significantly outperforms SPRank as well as other baseline methods. Among baselines, kNN-item is the best performing method in terms of P@5 and P@10 while BPRMF is the best performing baseline in terms of other evaluation metrics. A significant improvement of LODFM over BPRMF in MRR (+5.3%), MAP (+14.9%), nDCG@10 (+4.6%), P@10 (+12.9%) and R@10 (+8%) can be noticed. The results indicate that LOD-enabled features are able to improve the recommendation performance for factorization models. Compared to kNN-item, LODFM improves the performance by 8.2% and 2.4% in terms of P@5 and P@10, respectively. It is also interesting to observe that factorization models such as BPRMF and LODFM have much better performance especially in terms of recall compared to kNN-item. For example, LODFM improves the performance by 103%, 90% and 76.9% in terms of recall when $N = 1, 5$ and 10, respectively.

Table 3. Recommendation performance of LODFM using different sets of features such as property-object list (PO), subject-property list (SP) and PageRank scores (PR). The best performing strategy is in bold.

	PO	PO+SP	PO+PR	PO+SP+PR
MRR	0.5769	0.5403	0.5783	0.5561
MAP	0.2096	0.1957	0.2080	0.2008
nDCG@1	0.4224	0.3788	0.4236	0.3971
P@1	0.4224	0.3788	0.4236	0.3971
R@1	0.0237	0.0210	0.0241	0.0223
nDCG@5	0.4152	0.3861	0.4214	0.3963
P@5	0.3459	0.3222	0.3479	0.3280
R@5	0.0931	0.0841	0.0934	0.0866
nDCG@10	0.4904	0.4627	0.4945	0.4743
P@10	0.2973	0.2805	0.2975	0.2860
R@10	0.1558	0.1436	0.1541	0.1476

5.2 Model Analysis

Analysis of Features. To better understand the contributions of each feature set for recommendations, we discuss the recommendation performance with different sets of features for FM in this section. Table 3 shows the recommendation

performance of LODFM using different features with $m = 10$. The two fundamental features - user and item indexes are included by default and omitted from the table for clarity.

Overall, using a property-object list (PO) and the PageRank score (PR) of items provides the best performance compared to other strategies. As we can see from Table 3, PO+PR improves the recommendation performance compared to PO in terms of most of the evaluation metrics. Similar results can be observed by comparing PO+SP+PR against PO+SP, which shows the importance of PageRank scores of items. On the other hand, the performance is decreased by including SP, e.g., PO+SP vs. PO and PO+SP+PR vs. PO+PR. This shows that incoming knowledge about movie items is not helpful in improving recommendation performance.

Analysis of Dimensionality m for factorization. The dimensionality of factorization plays an important role in capturing pairwise interactions of input variables when m is chosen large enough [29]. Figure 4 illustrates the recommendation performance using different values for the dimensionality of factorization




Fig. 4. Recommendation performance based on different values for the dimensionality m of FM using PO+PR in terms of different evaluation metrics.

(The results of P@1 are equal to nDCG@1 and therefore omitted from Fig. 4(b)) using PO and PR as LOD-enabled features. As we can see from the figure, the performance consistently increases with higher values of m until $m = 200$ in terms of five evaluation metrics. For example, the performance is improved by 7.5% and 11.4% in terms of MRR and MAP with $m = 200$ compared to $m = 10$. There is no significant improvement with values higher than 200 for m .

6 Conclusions

In this paper, we investigated using FM with lightweight LOD-enabled features, such as property-object lists, subject-property lists, and PageRank scores of items which can be directly obtained from the DBpedia SPARQL Endpoint, for top- N recommendations. The results show that our proposed approach significantly outperforms compared approaches such as SPRank, BPRMF. In addition, we analyzed the recommendation performance based on different combinations of features. The results indicate that using the property-object list and the PageRank scores of items can provide the best performance. On the other hand, including the subject-property list of items is not helpful in improving the quality of recommendations but rather decreases the performance. In the future, we plan to evaluate our approach using other datasets in different domains. Furthermore, we aim to investigate other lightweight LOD-enabled features which might be useful to improve the recommendation performance.

Acknowledgments. This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289 (Insight Centre for Data Analytics).

References

1. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: A Nucleus for a Web of Open Data. In: Aberer, K., Choi, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) ASWC/ISWC -2007. LNCS, vol. 4825, pp. 722–735. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-76298-0_52](https://doi.org/10.1007/978-3-540-76298-0_52)
2. Di Noia, T., Cantador, I., Ostuni, V.C.: Linked Open Data-Enabled Recommender Systems: ESWC 2014 Challenge on Book Recommendation. In: Presutti, V., Stankovic, M., Cambria, E., Cantador, I., Di Iorio, A., Di Noia, T., Lange, C., Reforgiato Recupero, D., Tordai, A. (eds.) SemWebEval 2014. CCIS, vol. 475, pp. 129–143. Springer, Cham (2014). doi:[10.1007/978-3-319-12024-9_17](https://doi.org/10.1007/978-3-319-12024-9_17)
3. Di Noia, T., Mirizzi, R., Ostuni, V.C., Romito, D.: Exploiting the web of data in model-based recommender systems. In: Proceedings of the 6th ACM Conference on Recommender Systems, pp. 253–256. ACM (2012)
4. 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. ACM (2012)

5. 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, RecSys 2011, pp. 305–308. ACM, New York (2011)
6. de Gemmis, M., Lops, P., Musto, C., Narducci, F., Semeraro, G.: Semantics-Aware Content-Based Recommender Systems. In: Ricci, F., Rokach, L., Shapira, B. (eds.) Recommender Systems Handbook, pp. 119–159. Springer, Boston, MA (2015). doi:[10.1007/978-1-4899-7637-6_4](https://doi.org/10.1007/978-1-4899-7637-6_4)
7. Haveliwalla, T.H.: Topic-sensitive pagerank: a context-sensitive ranking algorithm for web search. *IEEE Trans. Knowl. Data Eng.* **15**(4), 784–796 (2003)
8. Heath, T., Bizer, C.: Linked data: evolving the web into a global data space. In: Synthesis lectures on the semantic web: theory and technology, vol. 1(1), pp. 1–136 (2011)
9. Heitmann, B.: An open framework for multi-source, cross-domain personalisation with semantic interest graphs. In: Proceedings of the sixth ACM conference on Recommender systems, pp. 313–316. ACM (2012)
10. Heitmann, B., Hayes, C.: Using linked data to build open, collaborative recommender systems. In: AAAI spring symposium: linked data meets artificial intelligence, pp. 76–81 (2010)
11. Heitmann, B., Hayes, C.: SemStim at the LOD-RecSys 2014 Challenge. In: Pre-sutti, V., Stankovic, M., Cambria, E., Cantador, I., Di Iorio, A., Di Noia, T., Lange, C., Reforgiato Recupero, D., Tordai, A. (eds.) SemWebEval 2014. CCIS, vol. 475, pp. 170–175. Springer, Cham (2014). doi:[10.1007/978-3-319-12024-9_22](https://doi.org/10.1007/978-3-319-12024-9_22)
12. Jeh, G., Widom, J.: SimRank: A measure of structural-context similarity. In: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2002, pp. 538–543. ACM, New York (2002)
13. Koren, Y.: Collaborative filtering with temporal dynamics. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2009, pp. 447–456. ACM, New York (2009)
14. Lalithsena, S., Kapanipathi, P., Sheth, A.: Harnessing relationships for domain-specific subgraph extraction: a recommendation use case. In: IEEE International Conference on Big Data, Washington D.C. (2016)
15. Leal, J.P.: Using proximity to compute semantic relatedness in RDF graphs. *Comput. Sci. Inf. Syst.* **10**(4), 1727–1746 (2013)
16. Lehmann, J., et al.: Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia. *Semant. Web J.* **6**(2015), 167–195 (2013)
17. Musto, C., Lops, P., Basile, P., de Gemmis, M., Semeraro, G.: Semantics-aware graph-based recommender systems exploiting linked open data. In: Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization, pp. 229–237. ACM (2016)
18. Musto, C., Narducci, F., Lops, P., De Gemmis, M., Semeraro, G.: ExplOD: A framework for explaining recommendations based on the linked open data cloud. In: Proceedings of the 10th ACM Conference on Recommender Systems, RecSys 2016, pp. 151–154. ACM, New York (2016)
19. Nguyen, P., Tomeo, P., Di Noia, T., Di Sciascio, E.: An evaluation of SimRank and personalized PageRank to build a recommender system for the web of data. In: Proceedings of the 24th International Conference on World Wide Web, pp. 1477–1482. ACM (2015)
20. Noia, T.D., Ostuni, V.C., Tomeo, P., Sciascio, E.D.: Sprank: semantic path-based ranking for top-n recommendations using linked open data. *ACM Trans. Intell. Syst. Technol. (TIST)* **8**(1), 9 (2016)

21. Oliveira, J., Delgado, C., Assaife, A.C.: A recommendation approach for consuming linked open data. *Expert Syst. Appl.* **72**, 407–420 (2017)
22. Ostuni, V.C., Di Noia, T., Di Sciascio, E., Mirizzi, R.: Top-n Recommendations from Implicit Feedback Leveraging Linked Open Data. In: *Proceedings of the 7th ACM Conference on Recommender Systems*, pp. 85–92. ACM (2013)
23. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank citation ranking: bringing order to the web. Technical report (1999)
24. Passant, A.: dbrec: Music Recommendations Using DBpedia. In: *ISWC 2010 SE - 14*, pp. 209–224 (2010)
25. Passant, A.: Measuring semantic distance on linking data and using it for resources recommendations. In: *AAAI Spring Symposium: Linked Data Meets Artificial Intelligence*, vol. 77, p. 123 (2010)
26. Piao, G., Ara, S., Breslin, J.G.: Computing the Semantic Similarity of Resources in DBpedia for Recommendation Purposes. In: Qi, G., Kozaki, K., Pan, J.Z., Yu, S. (eds.) *JIST 2015. LNCS*, vol. 9544, pp. 185–200. Springer, Cham (2016). doi:[10.1007/978-3-319-31676-5_13](https://doi.org/10.1007/978-3-319-31676-5_13)
27. Piao, G., Breslin, J.G.: Measuring semantic distance for linked open data-enabled recommender systems. In: *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, pp. 315–320. ACM (2016)
28. Rendle, S.: Factorization machines. In: *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pp. 995–1000. IEEE (2010)
29. Rendle, S.: Factorization machines with libFM. *ACM Trans. Intell. Syst. Technol.* **3**(3), 57:1–57:22 (2012)
30. 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, UAI 2009*, pp. 452–461. AUAI Press, Arlington (2009)
31. Rowe, M.: Transferring Semantic Categories with Vertex Kernels: Recommendations with SemanticSVD++. In: Mika, P., Tudorache, T., Bernstein, A., Welty, C., Knoblock, C., Vrandečić, D., Groth, P., Noy, N., Janowicz, K., Goble, C. (eds.) *ISWC 2014. LNCS*, vol. 8796, pp. 341–356. Springer, Cham (2014). doi:[10.1007/978-3-319-11964-9_22](https://doi.org/10.1007/978-3-319-11964-9_22)
32. Thalhhammer, A., Rettinger, A.: PageRank on Wikipedia: Towards General Importance Scores for Entities. In: Sack, H., Rizzo, G., Steinmetz, N., Mladenić, D., Auer, S., Lange, C. (eds.) *ESWC 2016. LNCS*, vol. 9989, pp. 227–240. Springer, Cham (2016). doi:[10.1007/978-3-319-47602-5_41](https://doi.org/10.1007/978-3-319-47602-5_41)
33. Wu, Q., Burges, C.J.C., Svore, K.M., Gao, J.: Adapting boosting for information retrieval measures. *Inf. Retrieval* **13**(3), 254–270 (2010)