

# KB4Rec: A Dataset for Linking Knowledge Bases with Recommender Systems

Wayne Xin Zhao, Gaole He, Hongjian Dou, Jin Huang, Siqui Ouyang and Ji-Rong Wen  
{batmanfly,ouyangsiqui0726}@gmail.com, {hegaole, hongjiandou, jin.huang, jrwen}@ruc.edu.cn  
School of Information, Renmin University of China

## Abstract

To develop a knowledge-aware recommender system, a key data problem is how we can obtain rich and structured knowledge information for recommender system (RS) items. Existing datasets or methods either use side information from original recommender systems (containing very few kinds of useful information) or utilize private knowledge base (KB). In this paper, we present a public linked KB dataset for recommender systems, named *KB4Rec v1.0*, which has linked three widely used RS datasets with the popular KB Freebase. Based on our linked dataset, we perform some interesting qualitative analysis experiments, in which we discuss the effect of two important factors (*i.e.*, popularity and recency) on whether a RS item can be linked to a KB entity.

## 1 Introduction

With the rapid development of Web techniques, various kinds of side information has become available in recommender systems (RS). In an early stage, such context information is usually unstructured, and its availability is limited to specific data domains or platforms [1, 2, 3]. Recently, more and more efforts have been made by both research and industry communities for structuring world knowledge or domain facts in a variety of data domains. One of the most typical organization forms is *knowledge base (KB)* [4]. KBs provide a general and unified way to organize and relate information entities, which have been shown to be useful in many applications. Specially, KBs have been

used in RSs [5, 6], usually called *knowledge-aware recommendation*.

To develop a knowledge-aware recommender system, a key data problem is how we can obtain rich and structured knowledge information for RS items. Overall, there are two main solutions from existing studies. First, side information is collected from the RS platform [1, 2, 3], and several studies further construct tiny and simple KB-like knowledge structure [7]. The number of attributes or relations is usually limited, and much useful knowledge information has not been considered. Second, several works propose to link RS with private KBs [5]. The linkage results are not publicly available.

To address the need for the linked dataset of RS and KBs, we present a public linked KB dataset for recommender systems, named *KB4Rec v1.0*, freely available at <https://github.com/RUCDM/KB4Rec>. Our basic idea is to heuristically link items from RSs with entities from a public large-scale KB<sup>1</sup>. On the RS side, we select three widely used datasets (*i.e.*, MovieLens [1], LFM-1b [2] and Amazon book [3]) covering three different data domains, namely movie, music and book; on the KB side, we select the well-known Freebase [8]. We try to maximize the applicability of our linked dataset by selecting very popular RS datasets and KBs. Specially, we are also aware of some closely related studies [9, 10], which also aim to link RS items with KB entities. While, our focus is on the Freebase, which is now widely used in many NLP or related domains [4].

In our KB4Rec v1.0 dataset, we organized the linkage results by linked ID pairs, which consists of a RS item ID and a KB entity ID. We do not share the original datasets, since they are maintained by original researchers or publishers. All the IDs are inner values from the original datasets. Once such a linkage has been accomplished, it is able to reuse existing

<sup>1</sup>We use the terms of “items” and “entities” respectively for RSs and KBs.

large-scale KB data for RSs. For example, the movie of “Avatar” from MovieLens dataset [1] has a corresponding entity entry in Freebase, and we are able to obtain its attribute information by reading out all its associated relation triples in Freebase. Based on the linked dataset, we first preform some interesting qualitative analysis experiments, in which we discuss the effect of two important factors (*i.e.*, popularity and recency) on whether a RS item can be linked to a KB entity. Finally, we present the comparison of several knowledge-aware recommendation algorithms on our linked dataset.

## 2 Existing Datasets and Methods

In this section, we briefly review the related datasets and methods.

Early knowledge-aware recommendation algorithms are also called context-aware recommendation algorithms, in which the side information from the original RS platform is considered as context data. For example, social network information of Epinions dataset is utilized in [11, 12], POI property information of Yelp dataset is utilized in [13], movie attribute information of MovieLens dataset is utilized in [7] and user profile information of microblogging dataset has been utilized in [14]. These datasets usually contain very few kinds of side information, and the relation between different kinds of side information is ignored.

To make such side information more structured, Heterogeneous Information Networks (HIN) have been proposed as a general technique for modeling information networks [15]. In HINs, we can effectively learn underlying relation patterns (called *meta-path*) and organize side information via meta-path-based representations. For example, HIN-based recommendation have been applied to solve PER [7] and MCRec [16]. HIN based algorithms usually rely on graph search algorithms, which is difficult to deal with large-scale relation pattern finding.

More recently, KBs have become a popular kind of data resources to store and organize world knowledge or domain facts. Many studies have been proposed [4] for the construction, inference and applications of KBs. Specially, several pioneering studies try to leverage existing KB information for improving the recommendation performance [17, 5, 18]. They apply a heuristic method for linking RS items with KB entities. In these studies, they use a private KB for linkage, which cannot be obtained publicly.

Specially, we are also aware of some closely related studies, including [9, 10], which also aim to link RS items with KB entities. While, our focus is on the Freebase, which is now widely used in many NLP or related domains [4].

## 3 Linked Dataset Construction

In our work, we need to prepare two kinds of datasets, namely RS and KB data. Next, we first give the detailed descriptions of the original datasets, and then discuss the linkage method.

**RS Datasets.** We consider three popular RS datasets for linkage, namely MovieLens, LFM-1b and Amazon book, which covers the three domains of movie, music and book respectively.

(1) *MovieLens* dataset [1] describes users’ preferences on movies. A preference record takes the form  $\langle \text{user}, \text{item}, \text{rating}, \text{timestamp} \rangle$ , indicating the rating score of a user for a movie at some time. There have been four MovieLens datasets released, known as *100K*, *1M*, *10M*, and *20M*, reflecting the approximate number of ratings in each dataset. We select the largest MovieLens *20M* for linkage.

(2) *LFM-1b* dataset [2] describes users’ interaction records on music. It provides information including artists, albums, tracks, and users, as well as individual listening events. It records the listening count of a song by a user, but does not contain rating information.

(3) *Amazon book* dataset [3] describes users’ preferences on book products with the data form of  $\langle \text{user}, \text{item}, \text{rating}, \text{timestamp} \rangle$ . The dataset is very sparse, containing 22 million ratings from 8 million users across nearly 23 million items.

In the three RS datasets, we several kinds of side information such as item titles (all), IMDB ID (movie), writer (book) and artist (music). We utilize such side information for subsequent KB linkage.

**KB Dataset.** We adopt the large-scale pubic KB *Freebase*. Freebase [8] is a KG announced by Metaweb Technologies, Inc. in 2007 and was acquired by Google Inc. on July 16, 2010. Freebase stores facts by triples of the form  $\langle \text{head}, \text{relation}, \text{tail} \rangle$ . Since Freebase shut down its services on August 31, 2016, we use the version of March 2015, which is its latest public version. We select Freebase because it has been widely applied in the research communities [4].

Table 1: Statistics of the linkage results. The three domains correspond to the RS datasets of MovieLens *20M*, LFM-1b and Amazon book, respectively.

Datasets	#Items	#Linked-Items	#Users	#Interactions
Movie	27,279	25,982	138,493	20,000,263
Music	6,479,700	1,254,923	120,317	1,021,931,544
Book	2,330,066	109,671	3,468,412	22,507,155

**RS to KB Linkage.** With an offline Freebase search API, we retrieve KB entities with item titles as queries. If no KB entity with the same title was returned, we say the RS item is *rejected* in the linkage process. If at least one KB entity with the same title was returned, we further incorporate one kind of side information as a refined constraint for accurate linkage: *IMDB ID*,

*artist name* and *writer name* are used for the three domains of movie, music and book respectively. We find only a small number (about one thousand for each domain) of RS items can not be accurately linked or rejected via the above procedure, and simply discard them. During the linkage process, we deal with several problems that will affect the results of string match algorithms, *e.g.*, lowercase, abbreviation, and the order of family/given names. Since the LFM-1b dataset is extremely large, we remove all the musics with fewer than ten listening events. Even after filtering, it still contains about 6.5 million musics.

**Basic Statistics.** We summarize the basic statistics of the three linked datasets in the second column of Table 1. It can be observed that for the MovieLens 20M dataset, we have a very high linkage ratio: about 95.2% items can be accurately linked to a KB entity. For LFM-1b dataset, the linkage ratio is 19.4%. But, the linkage ratio for the book domain is very low, about 4.7%. A possible explanation is that MovieLens 20M dataset contains fewer items than the other two datasets, which are ready refined by original releasers. Besides, we speculate that there may exist domain bias in the construction of Freebase. Although the linkage ratios for the latter two datasets are not high, the absolute numbers of linked items are large. Such a linked dataset is feasible for research-purpose studies.

**Shared Datasets.** We name the above linked linked KB dataset for recommender systems as *KB4Rec* v1.0, freely available at <https://github.com/RUCDM/KB4Rec>. In our KB4Rec v1.0 dataset, we organized the linkage results by linked ID pairs, which consists of a RS item ID and a KB entity ID. All the IDs are inner values from the original datasets. We have 25,982, 1,254,923, and 109,671 linked ID pairs for MovieLens 20M, LFM-1b and Amazon book respectively.

## 4 Linkage Analysis

Previously, we have shown the linkage ratios for different datasets. We find that a considerable amount of RS items can not be linked to KB entities. It is interesting to study what factors will affect the linkage ratio. We consider two kinds of factors for analysis.

**Effect of Popularity on Linkage.** Intuitively, a popular RS item should be more likely to be included in a KB than an unpopular item, since it is reasonable to incorporate more “important” RS items judged by the RS users into KBs. The construction of KB itself usually involves manual efforts, which is difficult to avoid the bias of human attention. To measure the popularity of a RS item, we adopt a simple frequency-based method by counting the number of users who have interacted with the item. This measure charac-

terizes the attractiveness of an item from the users in a RS. First, we sort the items ascendingly according to its popularity value. Then, we further equally divide all the items into five ordered bins with the same number of items. Hence, an item with a larger bin number will be more popular than another with a smaller bin number. Then we compute the linkage ratio for each bin and the results are reported in Fig. 1(a) (the three subfigures on the left). It can be observed that a bin with a larger number has a higher linkage ratio than the ones with a smaller number. The results indicate that popularity is likely to have positive effect on linkage.

**Effect of Recency on Linkage.** The second factor we consider is the recency, *i.e.*, the time when a RS item was created. Our assumption is that if a RS item was created or released on an earlier time, it would be more probable to be included in KBs. Since human attention aggregation is a gradually growing process, a RS item usually requires a considerable amount of time to become popular. To check this assumption, we need to obtain the release date of RS items. However, only the MovieLens 20M dataset contains such an attribute information, we only report the analysis result on this dataset. We first sort the items according to their release dates ascendingly, and then equally divide all the items into ten ordered bins following the procedure of the above popularity analysis. Finally, we compute the linkage ratios for each bin. The results are reported in Fig. 1(b). We can see that the linkage ratios gradually decrease with time going. The results indicate that recency is likely to have negative effect on linkage, *i.e.*, an older RS item seems to be more probable to be included in a KB than a more recent one. Especially, the last bin has a dramatic drop. A possible reason is that our dump of Freebase was released on March 2015, and many new items have not been included in Freebase.

## 5 Experiment

In this section, we present the comparison of some existing recommendation algorithms using our linked datasets.

**Experimental Setup.** Since our linked datasets are very large, we first generate a small test set for evaluation. We take the subset from the last year for LFM-1b dataset and the subset from year 2005 to 2015 for MovieLens 20M dataset. We also perform 3-core filtering for Amazon book dataset and 10-core filtering for other datasets. We consider the last-item recommendation task for evaluation. Since enumerating all the items as candidate is time-consuming, we pair each ground-truth with 100 negative items to form a ran-




Figure 1: Examining the effect of two factors on the linkage results. We use  $A, B, \dots$  to indicate the bin number in an ordered way. The first three subfigures correspond to the popularity analysis, and the last one corresponds to the recency analysis.

Table 2: Performance comparison of different methods on the task of last-item recommendation. We report the detailed statistics of the evaluation set in the first column.

Datasets	Methods	MRR	Hit@10	NDCG@10
<u>MovieLens 20M</u>	BPR	0.128	0.276	0.144
<i>#users:</i> 61,583	SVDFeature	0.204	0.448	0.243
<i>#items:</i> 19,533	mCKE	0.178	0.382	0.209
<i>#interactions:</i> 5,868,015	KSR	0.294	0.571	0.344
<u>LFM-1b</u>	BPR	0.227	0.458	0.265
<i>#users:</i> 7,694	SVDFeature	0.337	0.544	0.373
<i>#items:</i> 30,658	mCKE	0.371	0.541	0.399
<i>#interactions:</i> 203,975	KSR	0.427	0.607	0.460
<u>Amazon book</u>	BPR	0.222	0.505	0.272
<i>#users:</i> 65,125	SVDFeature	0.264	0.544	0.315
<i>#items:</i> 69,975	mCKE	0.248	0.494	0.291
<i>#interactions:</i> 828,560	KSR	0.353	0.653	0.413

domly candidate list. We adopt MRR, HR and NDCG as the evaluation metrics. To use the rich KB information, we embed KB data into low-dimensional vectors using extended KB subgraphs as in [6]. We consider four methods for comparisons, including **BPR** [19], **SVDFeature** [20], **mCKE** [5] and **KSR** [6]. BPR does not use KB information, SVDFeature and mCKE utilize KB information using shallow models, while KSR utilizes KB information using sequential neural networks and memory networks.

**Results and Analysis.** Table 2 presents the results of different methods for the last-item recommendation. First, among all the methods, BPR performs worst on three datasets, since it does not use KB information. Second, SVDFeature is better than BPR. It is implemented with a pairwise ranking loss function, and can be roughly understood as an enhanced BPR model with the incorporation of the learned KB embeddings. Finally, we analyze the performance of the knowledge-aware recommendation methods, namely mCKE and KSR. Overall, mCKE does not work well as expected, which only beats SVDFeature on the LFM-1b dataset. A possible reason is that our implementation of mCKE fixes the learned KB embeddings, while the original CKE model adaptively updates KB embeddings. As a comparison, the recently proposed KSR method works best consistently on the three datasets, which inte-

grates Recurrent Neural Networks with knowledge-enhanced Memory Networks for recommendation. The readers can refer to [6] for more detailed results and analysis.

## 6 Conclusion

This paper introduced a public dataset for linking RS with KB, namely *KB4Rec v1.0*. Our dataset covered three domains consists of a large number of linked ID pairs. As future work, we will consider linking more RS datasets with Freebase. We will also consider adopting other KB data for linkage, *e.g.*, YAGO and DBpedia.

## References

- [1] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets. *TiiS*, 5(4):1–19, 2016.
- [2] Markus Schedl. The lfm-1b dataset for music retrieval and recommendation. In *ICMR*, 2016.
- [3] Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*, 2016.
- [4] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE TKDE*, 2017.

- [5] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. Collaborative knowledge base embedding for recommender systems. In *SIGKDD*, pages 353–362, 2016.
- [6] Jin Huang, Wayne Xin Zhao, Hong-Jian Dou, Ji-Rong Wen, and Edward Y. Chang. Improving sequential recommendation with knowledge-enhanced memory networks. In *SIGIR*, 2018.
- [7] Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. Personalized entity recommendation: a heterogeneous information network approach. In *WSDM*, pages 283–292, 2014.
- [8] Google. Freebase data dumps. <https://developers.google.com/freebase/data>, 2016.
- [9] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. Dbpedia: A nucleus for a web of open data. In *ISWC 2007 + ASWC 2007*, 2007.
- [10] Tommaso Di Noia, Vito Claudio Ostuni, Paolo Tomeo, and Eugenio Di Sciascio. Sprank: Semantic path-based ranking for top- $N$  recommendations using linked open data. *ACM TIST*, 8(1):9:1–9:34, 2016.
- [11] Mohsen Jamali and Martin Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *RecSys*, pages 135–142, 2010.
- [12] Hao Ma, Irwin King, and Michael R. Lyu. Learning to recommend with social trust ensemble. In *SIGIR*, pages 203–210, 2009.
- [13] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. Content-aware point of interest recommendation on location-based social networks. In *AAAI*, pages 1721–1727, 2015.
- [14] Wayne Xin Zhao, Yanwei Guo, Yulan He, Han Jiang, Yuexin Wu, and Xiaoming Li. We know what you want to buy: a demographic-based system for product recommendation on microblogs. In *KDD*, 2014.
- [15] Yizhou Sun and Jiawei Han. Mining heterogeneous information networks: a structural analysis approach. *SIGKDD Explorations*, 14(2):20–28, 2012.
- [16] Binbin Hu, Chuan Shi, Wayne Xin Zhao, and Philip S. Yu. Leveraging meta-path based context for top- $N$  recommendation with a neural co-attention model. In *SIGKDD*, pages 1531–1540, 2018.
- [17] Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. DKN: deep knowledge-aware network for news recommendation. In *WWW*, pages 1835–1844, 2018.
- [18] Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. Ripple network: Propagating user preferences on the knowledge graph for recommender systems. 2018.
- [19] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, 2009.
- [20] Tianqi Chen, Weinan Zhang, Qiuxia Lu, Kailong Chen, Zhao Zheng, and Yong Yu. Svdfeature: a toolkit for feature-based collaborative filtering. *Journal of Machine Learning Research*, 2012.