

Bias in Knowledge Graphs – An Empirical Study with Movie Recommendation and Different Language Editions of DBpedia

Michael Matthias Voit ✉

University of Mannheim, Germany

Heiko Paulheim ✉ 

University of Mannheim, Germany

Abstract

Public knowledge graphs such as DBpedia and Wikidata have been recognized as interesting sources of background knowledge to build content-based recommender systems. They can be used to add information about the items to be recommended and links between those. While quite a few approaches for exploiting knowledge graphs have been proposed, most of them aim at optimizing the recommendation strategy while using a *fixed* knowledge graph. In this paper, we take a different approach, i.e., we fix the recommendation strategy and observe changes when using different underlying knowledge graphs. Particularly, we use different *language editions* of DBpedia. We show that the usage of different knowledge graphs does not only lead to differently biased recommender systems, but also to recommender systems that differ in performance for particular fields of recommendations.

2012 ACM Subject Classification Computing methodologies → Knowledge representation and reasoning; Information systems → Recommender systems

Keywords and phrases Knowledge Graph, DBpedia, Recommender Systems, Bias, Language Bias, RDF2vec

Digital Object Identifier 10.4230/OASICS.LDK.2021.14

Related Version *Full Version*: <https://arxiv.org/abs/2105.00674>

Supplementary Material *Software (Source Code)*: <https://github.com/voitjaner/Movie-RSs-Master-Thesis-Submission-Voit>

archived at swh.1.dir:5a1679a3579764cbd88c758be59c337e0f88a277

1 Introduction

Large-scale knowledge graphs, such as DBpedia [22] and Wikidata [41], are recognized as a valuable ingredient for intelligent applications [16]. In such applications, they can provide background information on the entities processed, which often leads to performance improvements in downstream processing steps [37].

In the past, different works have been proposed on building recommender systems based on knowledge graphs, most prominently, DBpedia. The first of those approaches has probably been *dbrec*, dating back to 2010 [32]. Since then, a number of approaches have been proposed, challenges around the topic have been conducted [7], and recent approaches have been utilizing the omnipresent knowledge graph embeddings for computing recommendations [29, 39].

The vast majority of those works always utilizes a fixed knowledge graph (DBpedia in most cases) and then optimizes the recommendation algorithm to provide the best empirical results on a test dataset. This means that by fixing the knowledge graph upfront, the influence of the chosen graph, its coverage, data quality, and possible biases, are not examined.



© Michael Matthias Voit and Heiko Paulheim;
licensed under Creative Commons License CC-BY 4.0
3rd Conference on Language, Data and Knowledge (LDK 2021).

Editors: Dagmar Gromann, Gilles Sérasset, Thierry Declerck, John P. McCrae, Jorge Gracia, Julia Bosque-Gil, Fernando Bobillo, and Barbara Heinisch; Article No. 14; pp. 14:1–14:13



OpenAccess Series in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper, we postulate that the choice of a particular knowledge graph has an influence on the behavior of the overall system, and may lead to a certain bias. To analyze this bias, we train a recommendation system with a fixed setup and parameter settings based on the embedding method RDF2vec [38], using different versions of DBpedia, which have been extracted from Wikipedia language editions.

Assuming that the coverage, quality, and level of detail of recommended items (in our test scenario: movies) varies from language edition to language edition, we expect a certain bias to shine up when using different knowledge graphs. This is confirmed by our experiments, however, the bias is not as obvious as we expected. While the straight forward assumption is that, e.g., a recommender system based on the German DBpedia edition would develop a stronger bias towards recommending German films, the effects are more subtle than that, exposing different significant biases with respect to production country and genre.

The rest of this paper is structured as follows. Section 2 discusses related works. In Section 3, we lay out our experiment set up, followed by an analysis of findings in Section 4. We conclude with a summary and an outlook on future work.

2 Related Work


The two most well-known families of recommender systems are collaborative filtering and content-based recommender systems. While the former analyze the behavior of users and recommends items that are consumed by users with a similar behavior as the one at hand, the latter exploit similarities between the items per se. For that category of approaches, a model of the recommended items is required, which can be unstructured (e.g., a textual description) or structured (e.g., a set of attributes). [35]

For structured representations, public knowledge graphs like DBpedia or Wikidata have been recognized as a valuable source of information, since they already contain a large amount of information on various items in a structured form [16]. Most classic approaches use DBpedia and/or knowledge graphs tailored to the domain at hand, and base their decision on a similarity function based on a set of hand-picked attributes (e.g., genre and artist for music; genre, director, and actor for movies).

The first generation of recommender systems based on Knowledge Graphs, such as dbrec, were based on hand-picking attributes and relations. Later approaches also exploited automatic approaches for selecting attributes, either by adapting measures such as TF-IDF to graph data [8], or by using machine learning methods such as Random Forests, which can be used on larger feature sets and automatically identify the relevant ones [36].

The most recent generation of such recommender system utilizes *knowledge graph embeddings* [11]. Such embedding methods project resources in a knowledge graph into a lower-dimensional, numerical vector space. Since many of those projection methods lead to vector spaces in which similar resources are close to each other, distance in the embedding space can be exploited for recommendation, as depicted in figure 1. Such approaches, among others, have been analyzed for RDF2vec [39], metapath2vec [52], TransE [6, 19, 20, 44, 49], TransR [25, 40, 44, 48, 53], TransH [4, 44], TransD [14, 44], ComplEx [19], LINE [48], Laplacian Eigenmaps and node2vec [29, 31], and embedding methods specifically tailored to the recommendation task, like *RippleNet* [43], *CFKG* [54], *Hierarchical Collaborative Embedding* [56], *MKR* [46], and *UPM* [57]. More recently, graph neural networks have also gained a bit of traction [45, 47].

While we are aware that this set of examples for the usage of knowledge graphs for content-based recommender systems is far from complete, a common trend can be observed in almost all publications about such systems: they always fix the knowledge graph to be



■ **Figure 1** 2-dimensional PCA projection of embedding vectors for a set of movies in DBpedia [39].

used upfront, and different variants are typically studied for the algorithms used to compute similarities, but not for the graph as such. In the rare cases where results obtained with multiple knowledge graphs are examined (e.g., [39], which contrasts results based on DBpedia and Wikidata), they are only compared based on the scoring function (e.g., F1 score), but other influences on the behavior of the recommender system are not analyzed. Hence, the influence of the choice for a particular knowledge graph is still underexplored.

In this paper, we conduct a study to shed some light on that aspect. To that end, we use different versions of DBpedia. DBpedia is extracted from Wikipedia infoboxes by the use of mappings to a central ontology. There are versions of DBpedia for different Wikipedia languages, called *DBpedia language editions* [22].

It is known that language editions of Wikipedia differ in coverage and level of detail. Their size ranges from a few hundred to a few million articles.¹ These differences have been analyzed with respect to various aspects, e.g., topical coverage [1, 9, 15, 28], article quality [24] and neutrality [3, 26, 51, 55], bias related to geography [2] or gender [42], and user behavior [12, 23].

The difference in the quality of infobox data in different Wikipedia language editions has also been studied [50]. This is particularly interesting for our scenario, since the DBpedia knowledge graph draws its information from those infoboxes. Hence, in the light of those studies, we expect significant differences in knowledge graphs extracted from different Wikipedia languages, and we want to explore in how far they lead to difference in downstream applications such as recommender systems.

¹ https://en.wikipedia.org/wiki/List_of_Wikipedias

14:4 Bias in Knowledge Graphs

■ **Table 1** Statistic of common movies in different language version with the movies mapped to the English DBpedia.

	it	pl	es	pt	fr	de	ru	nl	ja
# Movies	24k	13k	12k	12k	16k	19k	15k	10k	10k
Intersection with MovieLens 1M mapped to DBpedia-en	2,610	2,106	2,019	2,092	2,658	2,426	2,255	1,793	1,888

3 Experimental Setup

In our experiments, we use the MovieLens 1M dataset [13], which contains a 1M 1-5 star ratings by 6,040 users for 3,952 movies. Moreover, we use DBpedia version 2016-10.²

In earlier works, links from MovieLens 1M to DBpedia have been provided [30]. Since DBpedia has been evolving since the original linking was performed, we removed all links that refer either to entities which do not exist anymore (i.e., the URI has changed in later releases of DBpedia) or to entities derived from disambiguation pages. Our resulting dataset consists of 3,123 movies linked to the English DBpedia.

3.1 Datasets

To investigate the influence of the usage of different versions of DBpedia on a recommender system, we utilize different language versions of DBpedia. In a preliminary study, we looked at the ten largest language editions of DBpedia³, and analyzed the overlap with the 3,123 movies in our dataset linked to the English DBpedia. To that end, we utilize the links between DBpedia versions, which are extracted from the inter-language links in Wikipedia. The results are depicted in Table 1.

To ensure a reasonable coverage, we decided to use the five dataset which have the most information about movies and the highest overlap with the English dataset. Hence, we decided to base our analysis on English, Italian, French, German, and Russian. The subset of the original 3,123 movies which have a corresponding entity in all five datasets comprises 1,948 movies.

We apply two additional filtering steps, as suggested by [5], [30], and [38]. To avoid a popularity bias, the top-rated 1% of all movies are removed. In the second step, users with less than 50 ratings are removed, and so are movies without any ratings. After this step, we obtain a dataset with 1,918 movies, 675,960 ratings, and 3,642 users. This set is used as the basis for all our experiments.

3.2 Recommender Algorithm

As discussed above, in our experiments, we aim at keeping the recommender algorithm fixed, while varying the underlying knowledge graph. We intentionally use a simple algorithm for the recommendations, as our goal is not to maximize the performance of the recommendation as such, but to examine the influence of the underlying knowledge graph.

² <https://wiki.dbpedia.org/develop/datasets/dbpedia-version-2016-10>

³ <https://wiki.dbpedia.org/services-resources/datasets/dataset-statistics>

Following the setup in [38], we use RDF2vec to compute vector space embeddings of the different DBpedia graphs. RDF2vec extracts random walks from a knowledge graph, which are represented as “sentences” of entities and predicates in the knowledge graph. On that set of sentences, the word2vec algorithm [27] is run, which then computes an embedding vector for each entity (and predicate).

We computed RDF2vec embeddings for the five DBpedia language editions identified above, using the best performing parameter setting identified in [38], i.e., extracting 500 walks per entity with a depth of 4, a dimensionality of 200 for the word2vec model, using the Skip-Gram variant.⁴

The similarity of two movies is then computed as the cosine similarity of the corresponding vectors in that vector space. To that end, we compute a score y_{uj} for an unrated movie j and user u is calculated with the following formula:

$$y_{uj} = \frac{\sum_{i \in I_u} \cos(i, j) * r_{ui}}{\sum_{i \in I_u} \cos(i, j)} \quad (1)$$

Here, I_u are the previous observations from user u and r_{ui} denotes the rating of item i of the user u in the training set. Then, the N movies with the highest scores are returned for each user. For this procedure, we used the item similarity recommender of the GraphLab Create python framework⁵.

3.3 Metrics

To evaluate the quality of recommendations, we use the standard measures of recall, precision, and F1 score. Here, recall measures the fraction of items that a user ultimately rated positively which were recommended to him or her, and precision measures the fraction of recommended items which were ultimately rated positively. F1 is the harmonic mean between the two.

Besides the quality of recommendations, we are interested in differences among recommendations created based on different knowledge graphs. To that end, we look at different *categorical features*, like language or genre. For a categorical feature C (such as *language*), we can compute the probability of recommendations with a certain feature value c (such as *German*), i.e.,

$$p(c) = \frac{|R_c|}{|R|} \quad (2)$$

where $|R_c|$ is the total number of items that were recommended by a certain approach which have the categorical feature c , and $|R|$ is the total number of recommendations computed. These probabilities can then be compared for recommender systems based on different knowledge graphs.

In order to distinguish random variations from effects actually induced by the use of different knowledge graphs, we additionally conduct a chi-squared test:

$$\chi_{KG}^2 = \sum_{c \in C} \frac{(|R_c| - (c_e * |R|))^2}{(c_e * |R|)} \quad (3)$$

⁴ All code and data is available online at <https://github.com/voitjaner/Movie-RSs-Master-Thesis-Submission-Voit>

⁵ https://turi.com/products/create/docs/generated/graphlab.recommender.item_similarity_recommender.ItemSimilarityRecommender.html

■ **Table 2** Performance of the recommender systems per KG.

KG	Precision	Recall	F1 score
de	0.057	0.0404	0.047
fr	0.054	0.038	0.044
en	0.053	0.038	0.044
it	0.048	0.036	0.042
ru	0.042	0.028	0.034

where c_e is the expected value of recommendations with a categorical feature value c . We sum up the χ^2_{KG} values for all KGs, and compare them against the χ^2 distribution with $(|KGs| - 1) \cdot (|C| - 1)$ degrees of freedom, and an α value of 0.05 to test for significance. All the results presented below are not a random result according to that definition.

4 Findings and Observations

In total, we compare five recommender systems, based on the five different knowledge graphs. We analyze both the overall performance, as well as biases w.r.t. production countries and genres.

4.1 Overall Performance

In a first analysis, we look at the overall performance difference between the recommenders based on the five knowledge graphs. We can see that the one based on the German DBpedia works best, which is most likely due to a higher linkage degree for movies.

Most strikingly, the English DBpedia – which is used as the basis for the majority of works which claim to use “DBpedia” as a source of background knowledge – performs worse than its German and French counterpart. This shows that this choice, which is often done by simple heuristics such as popularity and availability, might not be an optimal one.

4.2 Bias for Production Countries

The first analysis for bias we perform is whether certain recommenders have stronger tendencies to recommend movies with a particular production country. The underlying hypothesis is that recommenders based on a knowledge graph derived from a Wikipedia in a particular language will have a tendency to also recommend more movies from a production country where that language is spoken (e.g., the recommender based on German DBpedia could have a stronger tendency to recommend German or Austrian movies).

Table 3 shows the top 10 production countries in the dataset. It can be observed that the dataset is heavily skewed towards movies from the USA and, to the lesser extent, UK, whereas other production countries only play a minor role.

Table 4 shows the fraction of movies from the top 10 production countries recommended by the systems based on the different knowledge graphs. We can see that the massive skew of the dataset towards US movies is also reflected in the results: except for the US movies, the majority of recommendations is below the expected value c_e .

Furthermore, it can be observed that, although significant differences in the behavior exist, there is no clear pattern of the form that follows the above mentioned hypothesis. Except for movies from the US and Australia, the peak of recommendations is always observed for a KG which is not in the respective language. For example, the fraction of German

■ **Table 3** Top 10 production countries in the dataset.

Country	# of movies	# of ratings
USA	1,679	622,946
UK	267	92,470
France	127	27,362
Germany	69	21,170
Italy	62	10,887
Canada	46	12,367
Australia	30	13,330
Japan	26	5,718
Spain	16	3,813
Mexico	15	6,790

■ **Table 4** Fraction of recommendations for different production countries by knowledge graph. c_e denotes the expected fraction based on the prevalence in the dataset.

Country/KG	de	fr	it	ru	en	c_e
USA	0.728	0.750	0.762	0.761	0.782	0.744
UK	0.136	0.143	0.098	0.091	0.108	0.110
France	0.028	0.030	0.036	0.037	0.026	0.033
Germany	0.012	0.018	0.012	0.030	0.034	0.025
Italy	0.016	0.009	0.013	0.009	0.009	0.013
Canada	0.020	0.009	0.021	0.005	0.006	0.015
Australia	0.017	0.010	0.013	0.008	0.020	0.016
Japan	0.006	0.005	0.012	0.004	0.006	0.007
Spain	0.006	0.004	0.006	0.002	0.005	0.005
Mexico	0.004	0.001	0.005	0.006	0.002	0.008

movies recommended by the system based on the English DBpedia is almost three times as high as the one based on the German DBpedia. Also, for other languages, the patterns are different: the highest fraction of French movies is recommended by the system based on the Russian KG, the highest fraction of Italian movies is recommended by the system based on the German KG, and so on.

4.3 Bias for Genres

In a second analysis, we inspect another possible bias induced by the different KGs, i.e., the bias to recommend movies from particular genres. Table 5 shows the top 10 genres in the dataset.

Table 6 shows the recommendations based on the different knowledge graphs for the top 10 genres. Here, we can again observe some interesting deviations. The recommender based on the Russian DBpedia has a tendency towards action, science fiction, and adventure movies, while the one based on the Italian DBpedia tends to recommend more movies from the comedy, thriller, and romance genres. Those findings partially correlate with studies on the popularity of particular genres in different countries. In [10], the authors discuss that, e.g., action movies are more popular in Russia than in English-speaking or European countries, and that comedy movies are more popular in Italy. Hence, it is likely that a local Wikipedia community in those countries will put more emphasis on editing movies in

■ **Table 5** Top 10 genres in the dataset.

Genre	# of movies	# of ratings
Drama	721	174,635
Comedy	562	184,700
Action	351	133,342
Thriller	320	74,457
Romance	244	44,784
Horror	225	22,700
Science Fiction	183	52,648
Adventure	172	36,827
Children's	130	10,316
Crime	115	7,621

■ **Table 6** Fraction of recommendations for different genres by knowledge graph. c_e denotes the expected fraction based on the prevalence in the dataset.

Genre/KG	de	fr	it	ru	en	c_e
Drama	0.198	0.170	0.187	0.172	0.190	0.162
Comedy	0.191	0.192	0.207	0.198	0.166	0.168
Action	0.089	0.010	0.074	0.129	0.112	0.123
Thriller	0.072	0.086	0.097	0.088	0.084	0.095
Romance	0.073	0.055	0.081	0.080	0.052	0.071
Horror	0.043	0.050	0.044	0.043	0.053	0.043
Science Fiction	0.055	0.045	0.044	0.056	0.053	0.073
Adventure	0.053	0.045	0.053	0.070	0.049	0.063
Children's	0.041	0.053	0.052	0.026	0.046	0.031
Crime	0.029	0.039	0.025	0.044	0.045	0.038

the respective genre in Wikipedia, which then leads those movies being more and better represented in the corresponding language-specific DBpedia, and ultimately a stronger bias of the recommender system based on that knowledge graph towards that genre.

4.4 Specific Performance Differences

The observation that recommenders based on different knowledge graphs expose biases towards particular genres also leads us to looking at the problem from a different angle. In particular, we want to analyze if recommenders based on different knowledge graphs work better or worse for single genres. To that end, we created partitions of our dataset by movie genre, and ran the recommender systems on those partitions. Overall, runs for ten different genres were performed.

The results are depicted in Table 7. We can observe that there are rather strong differences between the genre-specific recommender performance. The French DBpedia, which has also been identified as the best source of background knowledge above, yields superior results for half of the genres. On the other hand, also the Russian DBpedia, which shows the worst overall performance, outperforms all other recommender systems on the crime genre.

The differences on the individual genres are sometimes marginal, but for some genres (e.g., horror, children's), the best performing system can achieve results which are twice or even thrice as high as those which perform worst. This shows that there is no one-size-fits-all solution, and that the exploration of different knowledge graphs for a particular task and domain is at least as beneficial as the exploration of algorithmic alternatives.

■ **Table 7** Performance (F1) of recommenders for different genres by knowledge graph.

Genre/KG	de	fr	it	ru	en
Drama	0.040	0.045	0.034	0.030	0.040
Comedy	0.078	0.067	0.055	0.053	0.068
Action	0.091	0.114	0.089	0.080	0.105
Thriller	0.083	0.085	0.061	0.064	0.080
Romance	0.038	0.046	0.036	0.043	0.056
Horror	0.073	0.072	0.066	0.040	0.082
Science Fiction	0.101	0.124	0.106	0.080	0.095
Adventure	0.090	0.115	0.093	0.097	0.082
Children's	0.209	0.146	0.176	0.064	0.200
Crime	0.097	0.098	0.084	0.121	0.099

5 Conclusion and Future Work

In this paper, we have conducted a comparative study of recommender systems based on different knowledge graphs, particularly: versions of DBpedia, based on Wikipedia in different languages. The experiment design has been chosen in a way that a basic recommendation strategy was chosen and fixed, and five different underlying knowledge graphs were used. The results show that there are considerable differences in preferences of the recommenders. Particularly, we analyzed production countries and genres, but our method is generally applicable to other categorical variables as well (e.g., gender of producer or director, low, medium and high budget, etc.).

The second major observation is that despite overall trends, not all knowledge graphs are equally well suited for particular recommendation tasks. When building a recommender system for movies from a particular genre, the globally best performing knowledge graph might not be the one which performs best locally on a given task. Here, we argue that the choice of a knowledge graph – which is usually fixed upfront in most related works – should be treated equally, if not even more important as fine tuning algorithms.

The problem of fixing a knowledge graph upfront is not limited to recommender systems. Knowledge graphs have been suggested to be used in other fields as well, such as explainable AI [21], data interpretation [33], or social media analysis [34]. Like for recommender systems, biases induced by the choice of a particular knowledge graph have not been researched to a large extent here.

In the future, we see a few interesting directions to pursue. One of those is the extension of the analysis both to other domains, such as music or book recommendations, as well as the inclusion of further categorical variables, such as biases towards male or female authors, or black or white musicians.

The inclusion of further knowledge graphs in studies like these is also an interesting area. With the advent of more cross-domain knowledge graphs, such as Wikidata [41], CaLiGraph [17], and DBkWik [18] we assume that each of those comes with its very own coverage biases, and a setup like the one discussed in this paper would be a way of systematically investigating the possible impact of such biases on downstream applications. Furthermore, it is an open question whether combining information from different knowledge graphs is a suitable way of reducing the individual biases.

Finally, while we argue that the selection of a particular knowledge graph is at least as important as the selection and fine-tuning of a recommender algorithm, interaction effects between the two decisions must not be neglected. We assume that, while there is no one-size-

fits-all solution either on the knowledge graph nor on the algorithm side, the sweet spot for an optimal solution might not just be the straight forward combination of the knowledge graph and algorithm which perform best in isolation.

References

- 1 Patti Bao, Brent Hecht, Samuel Carton, Mahmood Quaderi, Michael Horn, and Darren Gergle. Omnipedia: bridging the wikipedia language gap. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1075–1084, 2012.
- 2 Pablo Beytía. The positioning matters: Estimating geographical bias in the multilingual record of biographies on wikipedia. In *Companion Proceedings of the Web Conference 2020*, pages 806–810, 2020.
- 3 Ewa S Callahan and Susan C Herring. Cultural bias in wikipedia content on famous persons. *Journal of the American society for information science and technology*, 62(10):1899–1915, 2011.
- 4 Yixin Cao, Xiang Wang, Xiangnan He, Zikun Hu, and Tat-Seng Chua. Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences. In *The world wide web conference*, pages 151–161, 2019.
- 5 Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 39–46, 2010.
- 6 Amine Dadoun, Raphaël Troncy, Olivier Ratier, and Riccardo Petitti. Location embeddings for next trip recommendation. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 896–903, 2019.
- 7 Tommaso Di Noia, Iván Cantador, and Vito Claudio Ostuni. Linked open data-enabled recommender systems: Eswc 2014 challenge on book recommendation. In *Semantic Web Evaluation Challenge*, pages 129–143. Springer, 2014.
- 8 Tommaso Di Noia, Roberto Mirizzi, Vito Claudio Ostuni, and Davide Romito. Exploiting the web of data in model-based recommender systems. In *Proceedings of the sixth ACM conference on Recommender systems*, pages 253–256, 2012.
- 9 Young-Ho Eom, Pablo Aragón, David Laniado, Andreas Kaltenbrunner, Sebastiano Vigna, and Dima L Shepelyansky. Interactions of cultures and top people of wikipedia from ranking of 24 language editions. *PloS one*, 10(3):e0114825, 2015.
- 10 Stephen Follows. The relative popularity of genres around the world:, 2016. URL: <https://stephenfollows.com/relative-popularity-of-genres-around-the-world>.
- 11 Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- 12 Noriko Hara, Pnina Shachaf, and Khe Foon Hew. Cross-cultural analysis of the wikipedia community. *Journal of the American Society for Information Science and Technology*, 61(10):2097–2108, 2010.
- 13 F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- 14 Ming He, Bo Wang, and Xiangkun Du. Hi2rec: Exploring knowledge in heterogeneous information for movie recommendation. *IEEE Access*, 7:30276–30284, 2019.
- 15 Brent Hecht and Darren Gergle. Measuring self-focus bias in community-maintained knowledge repositories. In *Proceedings of the fourth international conference on communities and technologies*, pages 11–20, 2009.
- 16 Nicolas Heist, Sven Hertling, Daniel Ringler, and Heiko Paulheim. *Knowledge Graphs on the Web – an Overview*, pages 3–22. IOS Press, 2020.
- 17 Nicolas Heist and Heiko Paulheim. Uncovering the semantics of wikipedia categories. In *International semantic web conference*, pages 219–236. Springer, 2019.

- 18 Sven Hertling and Heiko Paulheim. Dbkwik: A consolidated knowledge graph from thousands of wikis. In *2018 IEEE International Conference on Big Knowledge (ICBK)*, pages 17–24. IEEE, 2018.
- 19 Hen-Hsen Huang. An mpd player with expert knowledge-based single user music recommendation. In *IEEE/WIC/ACM International Conference on Web Intelligence-Companion Volume*, pages 318–321, 2019.
- 20 Jin Huang, Wayne Xin Zhao, Hongjian Dou, Ji-Rong Wen, and Edward Y Chang. Improving sequential recommendation with knowledge-enhanced memory networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 505–514, 2018.
- 21 Freddy Lecue. On the role of knowledge graphs in explainable ai. *Semantic Web*, 11(1):41–51, 2020.
- 22 Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morseay, Patrick Van Kleef, Sören Auer, et al. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web*, 6(2):167–195, 2015.
- 23 Florian Lemmerich, Diego Sáez-Trumper, Robert West, and Leila Zia. Why the world reads wikipedia: Beyond english speakers. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 618–626, 2019.
- 24 Włodzimierz Lewoniewski, Krzysztof Węcel, and Witold Abramowicz. Quality and importance of wikipedia articles in different languages. In *International Conference on Information and Software Technologies*, pages 613–624. Springer, 2016.
- 25 Qika Lin, Yaoqiang Niu, Yifan Zhu, Hao Lu, Keith Zvikomborero Mushonga, and Zhen-dong Niu. Heterogeneous knowledge-based attentive neural networks for short-term music recommendations. *IEEE Access*, 6:58990–59000, 2018.
- 26 Paolo Massa and Federico Scrinzi. Manypedia: Comparing language points of view of wikipedia communities. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*, pages 1–9, 2012.
- 27 Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- 28 Volodymyr Miz, Joëlle Hanna, Nicolas Aspert, Benjamin Ricaud, and Pierre Vanderghyest. What is trending on wikipedia? capturing trends and language biases across wikipedia editions. In *Companion Proceedings of the Web Conference 2020*, pages 794–801, 2020.
- 29 Cataldo Musto, Pierpaolo Basile, and Giovanni Semeraro. Embedding knowledge graphs for semantics-aware recommendations based on dbpedia. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 27–31, 2019.
- 30 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 Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1–34, 2016.
- 31 Enrico Palumbo, Giuseppe Rizzo, and Raphaël Troncy. Entity2rec: Learning user-item relatedness from knowledge graphs for top-n item recommendation. In *Proceedings of the eleventh ACM conference on recommender systems*, pages 32–36, 2017.
- 32 Alexandre Passant. dbrec – music recommendations using dbpedia. In *International Semantic Web Conference*, pages 209–224. Springer, 2010.
- 33 Heiko Paulheim. Generating possible interpretations for statistics from linked open data. In *Extended Semantic Web Conference*, pages 560–574. Springer, 2012.
- 34 Guangyuan Piao and John G Breslin. Exploring dynamics and semantics of user interests for user modeling on twitter for link recommendations. In *proceedings of the 12th international conference on semantic systems*, pages 81–88, 2016.
- 35 Francesco Ricci, Lior Rokach, and Bracha Shapira. Recommender systems: introduction and challenges. In *Recommender systems handbook*, pages 1–34. Springer, 2015.

- 36 Petar Ristoski, Eneldo Loza Mencía, and Heiko Paulheim. A hybrid multi-strategy recommender system using linked open data. In *Semantic Web Evaluation Challenge*, pages 150–156. Springer, 2014.
- 37 Petar Ristoski and Heiko Paulheim. Semantic web in data mining and knowledge discovery: A comprehensive survey. *Journal of Web Semantics*, 36:1–22, 2016.
- 38 Petar Ristoski, Jessica Rosati, Tommaso Di Noia, Renato De Leone, and Heiko Paulheim. Rdf2vec: Rdf graph embeddings and their applications. *Semantic Web*, 10(4):721–752, 2019.
- 39 Jessica Rosati, Petar Ristoski, Tommaso Di Noia, Renato de Leone, and Heiko Paulheim. Rdf graph embeddings for content-based recommender systems. In *CEUR workshop proceedings*, volume 1673, pages 23–30. RWTH, 2016.
- 40 Xiaoli Tang, Tengyun Wang, Haizhi Yang, and Hengjie Song. Akupm: Attention-enhanced knowledge-aware user preference model for recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1891–1899, 2019.
- 41 Denny Vrandečić and Markus Krötzsch. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85, 2014.
- 42 Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. It’s a man’s wikipedia? assessing gender inequality in an online encyclopedia. *arXiv preprint*, 2015. [arXiv:1501.06307](https://arxiv.org/abs/1501.06307).
- 43 Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. Ripplenet: Propagating user preferences on the knowledge graph for recommender systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 417–426, 2018.
- 44 Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. Dkn: Deep knowledge-aware network for news recommendation. In *Proceedings of the 2018 world wide web conference*, pages 1835–1844, 2018.
- 45 Hongwei Wang, Fuzheng Zhang, Mengdi Zhang, Jure Leskovec, Miao Zhao, Wenjie Li, and Zhongyuan Wang. Knowledge-aware graph neural networks with label smoothness regularization for recommender systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 968–977, 2019.
- 46 Hongwei Wang, Fuzheng Zhang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. Multi-task feature learning for knowledge graph enhanced recommendation. In *The World Wide Web Conference*, pages 2000–2010, 2019.
- 47 Hongwei Wang, Miao Zhao, Xing Xie, Wenjie Li, and Minyi Guo. Knowledge graph convolutional networks for recommender systems. In *The world wide web conference*, pages 3307–3313, 2019.
- 48 Meng Wang, Mengyue Liu, Jun Liu, Sen Wang, Guodong Long, and Buyue Qian. Safe medicine recommendation via medical knowledge graph embedding. *arXiv preprint*, 2017. [arXiv:1710.05980](https://arxiv.org/abs/1710.05980).
- 49 Xinyu Wang, Ying Zhang, Xiaoling Wang, and Jin Chen. A knowledge graph enhanced topic modeling approach for herb recommendation. In *International Conference on Database Systems for Advanced Applications*, pages 709–724. Springer, 2019.
- 50 Krzysztof Węcel and Włodzimierz Lewoniewski. Modelling the quality of attributes in wikipedia infoboxes. In *International Conference on Business Information Systems*, pages 308–320. Springer, 2015.
- 51 Hartmut Wessler, Christoph Kilian Theil, Heiner Stuckenschmidt, Angelika Storrer, and Marc Debus. Wikiganda: Detecting bias in multimodal wikipedia entries. *New Studies in Multimodality. London/New York: Bloomsbury*, pages 201–224, 2017.
- 52 Deqing Yang, Zikai Guo, Ziyi Wang, Juyang Jiang, Yanghua Xiao, and Wei Wang. A knowledge-enhanced deep recommendation framework incorporating gan-based models. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 1368–1373. IEEE, 2018.
- 53 Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. Collaborative knowledge base embedding for recommender systems. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 353–362, 2016.

- 54 Yongfeng Zhang, Qingyao Ai, Xu Chen, and Pengfei Wang. Learning over knowledge-base embeddings for recommendation. *Algorithms*, 11(9), 2018.
- 55 Yiwei Zhou, Elena Demidova, and Alexandra I Cristea. Who likes me more? analysing entity-centric language-specific bias in multilingual wikipedia. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, pages 750–757, 2016.
- 56 Zili Zhou, Shaowu Liu, Guandong Xu, Xing Xie, Jun Yin, Yidong Li, and Wu Zhang. Knowledge-based recommendation with hierarchical collaborative embedding. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 222–234. Springer, 2018.
- 57 Guiming Zhu, Chenzhong Bin, Tianlong Gu, Liang Chang, Yanpeng Sun, Wei Chen, and Zhonghao Jia. A neural user preference modeling framework for recommendation based on knowledge graph. In *Pacific Rim International Conference on Artificial Intelligence*, pages 176–189. Springer, 2019.