Knowledge Graph Datasets for Recommendation

Vincenzo Paparella[1,*], Alberto Carlo Maria Mancino[1,*], Antonio Ferrara[1], Claudio Pomo[1], Vito Walter Anelli[1] and Tommaso Di Noia[1]

[1]*Politecnico di Bari, Bari, Italy*

Abstract

In the era of daily information overload, personalized retrieval applications represent a crucial solution to provide suggestions to users. The research and industrial community have devoted an unprecedented effort to propose approaches and architectures to extract relevant and tailored information from every shred of knowledge. Inspired by the advances in knowledge graph, Graph Convolutional Networks, Link Prediction, and Recommender Systems research, this study aims to meet their cutting-edge research needs of holistic datasets by largely expanding the information available for two well-known recommendation datasets in the book and movie domains, i.e., LibraryThing and MovieLens 25M. We collect the associated knowledge graphs (KGs) for each of them and publish a mapping for each item in the two datasets with the corresponding entities in the original Wikidata, DBpedia, and Freebase KGs. Our work is available at https://github.com/sisinflab/Augmented-and-Linked-Open-Datasets-for-Recommendation.

Keywords

Recommender Systems, Dataset, Knowledge Graphs

1. Introduction and motivation

Recommender Systems (RSs) constitute the backbone of online platforms, pursuing the purpose of suggesting to customers novel products or services that might meet their interests and needs. The pivotal role of recommender systems [1] has been advantaged by the notable accuracy provided by collaborative filtering methods [2]. **Collaborative filtering (CF)** is the most traditional and still prominent recommendation approach. CF models calculate recommendations leveraging similarities in interaction/preference patterns of similar users. However, recommender systems generally have to face data sparsity issues due to the availability of little information about users' preferences, which often causes the so-called cold-start problem. To compensate for this lack of information, an increasing number of approaches combine collaborative information with auxiliary content attributes, such as tags, metadata, and geographical data. **Content-based Filtering (CBF)** models assemble a profile of user interests based on the description of the items that the user has previously consumed. Even though one could leverage this information using traditional CBF approaches like Vector Space Models and multimedia signals [3], knowledge graphs gave birth to specific research directions that comprise their own literature, models, and state-of-the-art. **Knowledge graphs (KGs)** have the advantage of covering a variety of heterogeneous domains, thus making it easier to transpose the advances in one domain to another. The adoption of knowledge graphs as a source of structured information has generated several advancements in the tasks strictly related to recommendation [4]: knowledge completion [5], preference elicitation [6], and user modeling [7, 8]. Given the graph nature of data, several different recommendation approaches have been proposed to tackle: knowledge transfer across domains [9], interpretable/explainable-recommendation [10, 11, 12], user modeling [7, 13, 14], and graph-based recommendation [15, 16, 17, 18, 19]. Although the KGs are all different, every method needs two side information: a mapping that establishes a connection between items within the catalog and entities in the KGs, and a set of RDF triples that utilize this mapping to describe the relevant domain.

Unfortunately, dealing with these information sources raises two important issues for researchers. First, it is usually a time and computationally expensive task due to the limits and constraints of KG services that make this operation frustrating. For instance, the public DBpedia Virtuoso SPARQL endpoint sets limits both on the rate of requests (whose violation could lead to a permanent ban) and on the maximum SPARQL query solution size[1]. Second, the information available on the Web rapidly and constantly changes over time. This causes the different retrieval operations of item attributes to be inconsistent and lead to different outcomes in recommendation research, thus dangerously jeopardizing the reproducibility of experiments [20, 21].

This study aims to tackle these issues by releasing an

Fifth Knowledge-aware and Conversational Recommender Systems (KaRS) Workshop @ RecSys 2023, September 18–22 2023, Singapore.

*Corresponding authors.

vincenzo.paparella@poliba.it (V. Paparella); alberto.mancino@poliba.it (A. C. M. Mancino); antonio.ferrara@poliba.it (A. Ferrara); claudio.pomo@poliba.it (C. Pomo); vitowalter.anelli@poliba.it (V. W. Anelli); tommaso.dinoia@poliba.it (T. Di Noia)

[1]https://www.dbpedia.org/blog/new-dbpedia-usage-report-2021/

extensive collection of side information exploited in tasks involving KGs. We focus on two well-known recommendation datasets in the community and detail the steps to gather the collection. In detail, the contribution of the work at hand is manifold:

- **a mapping between the items** in the LibraryThing (LT) and MovieLens 25M (ML25M) catalogs **and the entities available in three well-known knowledge graphs**: Wikidata [22], DBpedia [23], and Freebase [24];
- **the retrieval of DBpedia and Wikidata KG's triples** up to two hops to collect structured information connected to these resources, thus providing persistent and ready-to-use enriched datasets for performing reproducible experiments;
- **a methodology for datasets' items linking** as an approach to promote the emergence of other state-of-the-art enriched datasets;
- **an evaluation of the linking methodology** by comparing it with a state-of-the-art entity linker showing that traditional linking approaches are ineffective for this task.

The remainder of this paper is organized as follows. Section 2 surveys the existing dataset and the knowledge-aware recommendation works that use them. In Sections 3 and 4, we describe the outcome of our work, along with details about challenges we have faced and methodologies we have adopted. In Section 5, we evaluate the adopted text similarity-based solution for our work. Finally, Section 6 concludes this paper, proposing and discussing future directions.

2. Knowledge Graphs in Recommendation Works

According to Sun et al. [25], the ten most used datasets in the recent literature (sorted in descending order of adoption) are ML-1M [26], Netflix, LastFM, Yelp, ML-20M, Epinions [27], ML-10M, ML-100K, Book-Crossing [28], Amazon Music. Unsurprisingly, they relate with streaming and e-commerce platforms and cover the typical domains of movies, music, books, and shops. When linked data is available, it can be leveraged and adopted for addressing several tasks. For instance, Piao and Breslin [29] link MovieLens 1M and DBook datasets to the DBpedia KG and investigate the idea of transfer learning between item recommendations and KG completion tasks leveraging a co-factorization model. Also Vagliano et al. [30] and Alshammari and Nasraoui [31] use knowledge-aware versions of MovieLens for powering their models. He et al. [5] adopt in their model the public linking KB4Rec [32] and exploit implicit entity preferences to infer the plausibility of KG facts. Wang et al. [33] explore the intents behind a user-item interaction by using auxiliary information in Amazon Books, LastFM and Alibaba-iFashion [34]. Sun et al. [35] present an approach for KG embedding able to model and infer various relation patterns. They evaluate the proposed model on four widely used KGs, i.e., FB15k [36], WN18 [36], FB15k-237, and WN18RR. Zhang et al. [37] use MovieLens 1M and IntentBooks [38] to evaluate CKE, a knowledge-aware model extracting feature representations from a knowledge base and captures relationships between users and items. This large amount of work using side information and enriched datasets shows the potential of these tools. However, often these works use datasets whose linking strategies are highly different from each other and lack details for their reproducibility, putting reuse within other works at risk.

For instance, concerning MovieLens and LibraryThing, which we treat in this paper, few significant attempts exist to link these datasets to KGs. Wang et al. [39] link MovieLens 1M with IMDb information, such as genres, actors, directors, and writers to use it in a knowledge-aware model. Ostuni et al. [40] publish a mapping of MovieLens 1M to the DBpedia KG, while Zhao et al. [41] propose KB4Rec, a public linking of MovieLens 20M, LastFM 1B, and Amazon Book to the popular KB Freebase. Moreover, to the best of our knowledge, only Di Noia et al. [42] in 2016 succeded linking LibraryThing to a KG, specifically to DBpedia, likely due to some domain limitations detailed in Section 4.

These efforts are often not complemented by details about the adopted linking methodologies and do not provide an evaluation of their outcome. In addition, previous literature provides linked datasets but does not publicly release the RDF triples retrievable by exploring the KGs. This requires the researchers to scan the KGs each time, thus causing high bandwidth, resource, and time consumption. It is worth highlighting that *entities and triples of the KGs change over time, thus producing a mismatch of the retrieved triples that causes* ***reproducibility issues***.

These drawbacks have animated us to release extensive linkings of the large MovieLens 25M and LibraryThing datasets to DBpedia, Wikidata, and Freebase and describe our methodology in detail. We provide researchers with a vast, up-to-date, and consistent-over-time collection of RDF triples. The proposed methodology permits the reproducibility of the results and helps researchers in extracting novel datasets for other domains. Finally, we propose an evaluation of our linking strategy, thus enhancing the trust of the research community in using the brand-new enriched datasets.

Table 1
MovieLens 25M and LibraryThing datasets statistics referred to user-item registered interactions.

	MovieLens 25M	LibraryThing
Ratings	25 000 095	1 707 070
Items	59 047	506 165
Users	162 541	83 194
Sparsity	$0,997$	$0,999$

3. Dataset Enrichment

One of the main contributions of this work is providing the community with two enriched versions of MovieLens 25M (ML25M) and LibraryThing (LT), thus guaranteeing the reproducibility of any experiment making use of additional knowledge. Table 1 reports the the overall statistics for the original datasets, while Table 2 summarizes the collected resources for both datasets. Each resource is extensively detailed in the **GitHub repository**.

MovieLens 25M. The MovieLens datasets are the result of a collective effort of thousands of people expressing explicit preferences about movies [26]. These datasets constantly increase in size regarding ratings, users, and items, with the release that we hereby consider including about 25 million ratings. We have complemented the items in MovieLens 25M with links to resources of Wikidata, DBpedia, and Freebase knowledge graphs. With this information available, we have explored the Wikidata and DBpedia graphs up to the second hop [43, 44] by querying the public SPARQL endpoints to retrieve the subgraphs originating from the entities associated with the dataset items.

On the one hand, item linkings are collected comprising the item ID and the URIs referring to the same entity in the three knowledge bases. On the other hand, the facts retrieved from the knowledge graphs have been provided in the standard Resource Description Framework (RDF). RDF provides a simple data model to encode knowledge in a structured way through a set of triples in the form $\langle \sigma, \rho, \omega \rangle$. Each triple represents the connection $\sigma \xrightarrow{\rho} \omega$ between the subject σ and the object ω through a predicate ρ. We define **exploration at 1-hop** of $\overline{\sigma}$ as the operation of retrieving the triples having $\overline{\sigma}$ as the subject. Generalizing, the **exploration at n-hop** is the operation to retrieve the triples having as subjects all the objects obtained at the exploration from 1-hop to $(n-1)$-hops. In the proposed resources, we collected the triples obtained exploring DBpedia and Wikidata at 1-hop and at 2-hops.

LibraryThing. The LT dataset has been collected in 2011 from the namesake website, which allows users to create a virtual catalog of the books they own or have read. The dataset contains item reviews with ratings from 1 to 5 along with the number of "helpfulness" votes they received. The linking of LibraryThing items to DBpedia and Wikidata URIs followed a similarity-based approach, due to some information mismatching that we discuss in Section 4 and that made difficult to obtain a definite mapping. As a consequence, the item-URI linking comes with a similarity, that accounts for a score that a book actually corresponds to the associated resource on Wikidata and DBpedia. Hence, one can set a confidence threshold, bearing in mind that the lower is the threshold, the higher is the number of linked entities, but the higher is the occurrence of false positives. Although a threshold in the range of $[0.5, 0.6]$ may represent an acceptable trade-off between true positives and false positive, we have not filtered the retrieved information, leaving each researcher the choice to select it. Regardless of the linking similarity scores, we have linked the items to Freebase entities. Additionally, we have provided all the triples retrievable at July 2023 from DBpedia and Wikidata by querying the SPARQL endpoints using the linked entities and exploring the graphs at the first and at the second hop.

4. Item Linking Methodology

In this section, we present methodologies, challenges, and solutions adopted for performing the item linkings described in Section 3. Our goal is to sketch best practices and highlight issues that raise from item linking, with the aim of encouraging the discussion about the need of a common workflow to collect augumented datasets.

The goal of an item linker is to connect the items of a catalog with the corresponding resources in the Linked Open Data (LOD) Knowledge Base. With this step, we try to fill the gap between the items in the catalog and LOD cloud. In this respect, we could adopt two different solutions to address this goal, i.e., *direct item linking* and *item description linking*, as suggested by Di Noia [45].

MovieLens 25M. MovieLens 25M is provided with a map from internal item IDs to the identifiers of the movies in the massive Internet Movie Database (IMDb) and in The Movie Database (TMDb). We have exploited this connection and adopted a *direct item linking* strategy, querying the Wikidata endpoint with a SPARQL query. Specifically, we retrieve each item resource URI in MovieLens 25M by performing at least one SPARQL query: *(i)* the first query uses the IMDb identifier to find the corresponding entity on Wikidata; *(ii)* if the first attempt does not succeed, another query is sent to the endpoint using the TMDb identifier. Once the resource has been found in Wikidata, we reconstruct the corresponding

Table 2
Statistics of the collected resources categorized by related datasets and data sources.

Dataset	Source	Linked	Knowledge Graph Collection					
			Hop	*Retrieved triples*	*Subject*	*Objects*	*Predicates*	*Graph density*
ML25M	**Wikidata**	54 563	*1st*	4 363 390	54 298	3 212 340	1 174	0.041×10^{-5}
			2nd	18 481 584	3 095 341	7 545 673	7 639	0.018×10^{-5}
	Dbpedia	48 948	*1st*	6 365 010	48 822	1 284 078	268	0.370×10^{-5}
			2nd	51 517 943	524 673	12 401 997	5 990	0.033×10^{-5}
	Freebase	43 612						
LT	**Wikidata**	65 740	*1st*	1 537627	91 848	1 129417	1 443	0.105×10^{-5}
			2nd	8 333 024	987 788	4 654 913	7 763	0.029×10^{-5}
	Dbpedia	53 088	*1st*	5 098 254	72 066	1 505 610	1 038	0.214×10^{-5}
			2nd	80 305 409	809 001	17 288 244	6 778	0.026×10^{-5}
	Freebase	54 006						

DBpedia URI from the English Wikipedia page URL. The linking approach chosen for MovieLens 25M does not generate false positives thanks to the accurate references from MovieLens 25M items to IMDBb and TMDb IDs, and to the explicit reference to those IDs in the Wikidata Knowledge Graph. However, for each item in the dataset, at least one SPARQL query had to be performed. Therefore, *the researcher should take into account the costs for establishing a huge number of HTTP requests and the limitations that an endpoint may place.*

LibraryThing. LibraryThing collects books using internal identifiers without sharing the international identifier or a website page. Moreover, books are usually sold in different editions and LibraryThing manages this information treating them as completely different books. Although book entities in Wikidata and DBpedia are generally linked to a page of the LibraryThing website, a relevant issue comes from the fact that the LibraryThing dataset and the Wikipedia/DBpedia KGs often link to different editions of the same book — thus generating a mismatching of the identifiers. To overcome this issue, we adopt the *item description linking* strategy. With this approach, we exploit a similarity function to compare a resource's title and author available in the dataset with the corresponding information of the books from Wikidata or DBpedia. Since the similarity is computed with a dedicated offline operation, this approach requires to query only once Wikidata and DBpedia to retrieve all the books in these KGs along with their titles and authors. The linking task is then realized with a post-hoc analysis. To summarize, we created the collection B of the books' URI instances available on Wikidata and DBpedia with their titles (B_t) and authors (B_a). Then, we created the collection LT of books IDs from LibraryThing with their titles (LT_t) and authors (LT_a). Finally, given an instance $b \in B$ with its title $b_t \in B_t$ and its author $b_a \in B_a$, we aim to create a direct mapping $b \rightarrow lt$, where $lt \in LT$ is characterized by a title $lt_t \in LT_t$ and an author $lt_a \in LT_a$, based upon the similarities between titles labels b_t and lt_t and authors labels b_a and lt_a.

To this end, we consider and combine three popular similarity metrics, namely the Jaro [46], Ratio [47], and n-grams [48] similarities. We assume string similarity could be a reasonable solution to ensure that two entities are the same. Furthermore, by calculating and combining the similarities between title names and authors, we may address critical situations where titles are the same, but authors do not match, and vice versa. Then, by considering two strings s_1 and s_2, we define $m\_sim(s_1, s_2)$ as the average of the these similarity measures. Concerning the B and LT sets, and taking into account an instance $b \in B$ and an instance $lt \in LT$, we perform $m\_sim(b_t, lt_t)$ (i.e., the similarity between title labels) and $m\_sim(b_a, lt_a)$ (i.e., the similarity between author labels). Subsequently, we multiply them to obtain the combined similarity measure $c_{sim}(b, lt) = m\_sim(b_t, lt_t) * m\_sim(b_a, lt_a)$. We set $m\_sim(b_a, lt_a) = 1$ if the labels of the author b_a and/or lt_a are missing, so that the title labels are the only ones to contribute to the similarity since 1 is the neutral factor for the multiplication. We perform this operation for each instance $lt \in LT$, obtaining a number of triples $t \triangleq (b, lt, c_{sim}(b, lt))$ equals to $|LT| \times |B|$. Finally, we select at most one triple t with $\max(c_{sim}(b, lt))$ for each combination of b and lt. In this respect, we obtain a direct and unambiguous mapping between the book ID of b and the URI in lt. However, the linking may not be unique at this stage because some different URIs may be associated with the same ID. In such cases, we make the linking unique by only considering the one characterized by the highest similarity value. The number of false positives when measuring the string similarity is strictly dependent on the threshold value chosen (see the discussion in Section 3).

From DBpedia and Wikidata to Freebase. Both the DBpedia and Wikidata KG provide between their entities and Freebase KG. Therefore, we get from the both KGs all the object entities referenced by these predicates to accomplish this linking stage. As for MovieLens 25M, a Freebase reference was searched by exploring DBpedia, while for the entities not linked to Freebase yet, a reference was searched by querying Wikidata. Concerning LibraryThing, the solution adopted to perform the item linking led us to generate two different comprehensive resources, one for each KG.

5. Evaluation

This section describes the limitations of existing named entity linkers [49, 50, 51] for this linking task and the need for an ad-hoc linker for the book domain. Then, an experimental evaluation is conducted to assess if the proposed solution can overcome the performance of existing linkers. In detail, differently from a traditional linking task where a textual context is provided, only titles and authors are retrieved from LibraryThing. Although this could already be a reasonable motivation, we compare the solution with a state-of-the-art linker, DBpedia Spotlight, a framework for automatic annotation of DBpedia entities from natural language text [50, 52].

However, a ground truth is needed to evaluate the methods. Unfortunately, LibraryThing does not provide this information, but MovieLens does. In fact, the MovieLens dataset contains titles, IMDb, and TMDb IDs. Therefore, we can exploit it to evaluate the performance of the linkers by fixing the available side information, i.e., the label of the film and its director taken from IMDb.

On the one hand, to get DBpedia Spotlight working at its best, we supply this framework with a context by generating sentences exploiting the available side information. For instance, given as metadata the film title "*Toy Story*" and its director "*John Lasseter*", we ask DBpedia spotlight to annotate DBpedia entities in the sentence "*The film Toy Story is directed by John Lasseter*". Then, DBpedia Spotlight should be capable of taking out, among the other recognized entities, the resource URI in the DBpedia Knowledge Graph (KG) of the film "*Toy Story*", which should correspond to the one found with direct item-linking. On the other hand, we follow the pipeline proposed for the book domain (see Section 4), retrieving all the films in the DBpedia KG with their labels and directors (corresponding to the book titles and their authors in the book domain) and computing string similarities with IMDb-derived data (the candidate items in this evaluation). Finally, we link a URI resource of DBpedia to the film having the highest similarity score and compare this linking with the ground truth. It is worth mentioning that side information could be diverse even though it refers to the same subject. For clarity, one can easily understand that "*John Lasseter*" and "*J. Lasseter*" are correct information about *Toy Story*'s director. However, even though this situation is common, an automated agent finds it more challenging to identify the same entity. Hence, we repeated the experiment to evaluate the robustness of the linkers by introducing a random noise on the strings. Specifically, each character is deleted e/o perturbed with a probability of 0.15.

Table 3
Comparison of successful item linking performed by DBpedia Spotlight and our linker for the movie domain.

Linked Items of ML 25M dataset	**Original Side Information**	**Noised Side Information**
DBpedia Spotlight	15 483	1 755
Ours	40 189	35 639

The results in Table 3 show the efficacy of the proposed solution. With the support of the original side information only, our linker successfully links 40 189 of the items, an acceptable performance compared to DBpedia Spotlight (15 483). For what concerns the robustness analysis, DBpedia Spotlight performance collapses up to 1755 of correct links, while the proposed linker loses only 4550 links. In conclusion, the results justify the adoption of a tailored linker for this peculiar application scenario.

6. Conclusion and Future Work

This work provides supplementary data of two well-known and widely used recommendation datasets (i.e., MovieLens 25M and LibraryThing) that can be utilized in tasks involving Knowledge Graphs. Inspired by the new information needs of cutting-edge research, we provide mappings to three well-known Knowledge Graphs: DBpedia, Wikidata, and Freebase. To ensure reproducibility of the future experiments, we also provide domain Knowledge Graphs at the first and second hop of exploration in the form of RDF triples, thus enabling reliable Graph-based and Knowledge-aware recommendation research. Furthermore, they are aligned and expose the same additional information to ease the researcher's effort in conducting evaluation and comparisons. For future work, we encourage the enrichment of datasets from other domains such as music and shops and their integration with spectrograms and images, respectively, to design novel multi-modal recommendation methods.

References

[1] A. C. M. Mancino, T. D. Noia, Towards differentially private machine learning models and their robustness to adversaries, in: ICWE, volume 13362 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 455–461.

[2] V. Paparella, D. D. Palma, V. W. Anelli, T. D. Noia, Broadening the scope: Evaluating the potential of recommender systems beyond prioritizing accuracy, in: RecSys '23: Seventeenth ACM Conference on Recommender Systems, Singapore, Singapore, 18 September 2023 - 22 September 2023, ACM, 2023. URL: https://doi.org/10.1145/3604915.3610649. doi:10.1145/3604915.3610649.

[3] F. A. Merra, V. W. Anelli, T. D. Noia, D. Malitesta, A. C. M. Mancino, Denoise to protect: A method to robustify visual recommenders from adversaries, in: SIGIR, ACM, 2023, pp. 1924–1928.

[4] V. W. Anelli, T. Di Noia, 2nd workshop on knowledge-aware and conversational recommender systems - kars, in: CIKM, ACM, 2019, pp. 3001–3002.

[5] G. He, J. Li, W. X. Zhao, P. Liu, J. Wen, Mining implicit entity preference from user-item interaction data for knowledge graph completion via adversarial learning, in: Y. Huang, I. King, T. Liu, M. van Steen (Eds.), WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020, ACM / IW3C2, 2020, pp. 740–751. URL: https://doi.org/10.1145/3366423.3380155. doi:10.1145/3366423.3380155.

[6] V. W. Anelli, R. D. Leone, T. D. Noia, T. Lukasiewicz, J. Rosati, Combining RDF and SPARQL with cp-theories to reason about preferences in a linked data setting, Semantic Web 11 (2020) 391–419. URL: https://doi.org/10.3233/SW-180339. doi:10.3233/SW-180339.

[7] H. Wang, F. Zhang, J. Wang, M. Zhao, W. Li, X. Xie, M. Guo, Exploring high-order user preference on the knowledge graph for recommender systems, ACM Trans. Inf. Syst. 37 (2019) 32:1–32:26. URL: https://doi.org/10.1145/3312738. doi:10.1145/3312738.

[8] V. W. Anelli, T. D. Noia, E. D. Sciascio, A. Ferrara, A. C. M. Mancino, Sparse embeddings for recommender systems with knowledge graphs, in: IIR, volume 2947 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021.

[9] Q. Zhang, P. Hao, J. Lu, G. Zhang, Cross-domain recommendation with semantic correlation in tagging systems, in: International Joint Conference on Neural Networks, IJCNN 2019 Budapest, Hungary, July 14-19, 2019, IEEE, 2019, pp. 1–8. URL: https://doi.org/10.1109/IJCNN.2019.8852049. doi:10.1109/IJCNN.2019.8852049.

[10] V. W. Anelli, V. Bellini, T. D. Noia, E. D. Sciascio, Knowledge-aware interpretable recommender systems, in: I. Tiddi, F. Lécué, P. Hitzler (Eds.), Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges, volume 47 of *Studies on the Semantic Web*, IOS Press, 2020, pp. 101–124. URL: https://doi.org/10.3233/SSW200014. doi:10.3233/SSW200014.

[11] V. W. Anelli, T. D. Noia, E. D. Sciascio, A. Ragone, J. Trotta, Semantic interpretation of top-n recommendations, IEEE Trans. Knowl. Data Eng. 34 (2022) 2416–2428. URL: https://doi.org/10.1109/TKDE.2020.3010215. doi:10.1109/TKDE.2020.3010215.

[12] Z. Yang, S. Dong, Hagerec: Hierarchical attention graph convolutional network incorporating knowledge graph for explainable recommendation, Knowl. Based Syst. 204 (2020) 106194. URL: https://doi.org/10.1016/j.knosys.2020.106194. doi:10.1016/j.knosys.2020.106194.

[13] R. Ojino, User's profile ontology-based semantic model for personalized hotel room recommendation in the web of things: student research abstract, in: C. Hung, G. A. Papadopoulos (Eds.), Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019, Limassol, Cyprus, April 8-12, 2019, ACM, 2019, pp. 2314–2316. URL: https://doi.org/10.1145/3297280.3297661. doi:10.1145/3297280.3297661.

[14] A. Ferrara, V. W. Anelli, A. C. M. Mancino, T. D. Noia, E. D. Sciascio, Kgflex: Efficient recommendation with sparse feature factorization and knowledge graphs, ACM Trans. Recomm. Syst. (2023). URL: https://doi.org/10.1145/3588901. doi:10.1145/3588901, just Accepted.

[15] V. W. Anelli, Y. Deldjoo, T. D. Noia, D. Malitesta, V. Paparella, C. Pomo, Auditing consumer- and producer-fairness in graph collaborative filtering, in: J. Kamps, L. Goeuriot, F. Crestani, M. Maistro, H. Joho, B. Davis, C. Gurrin, U. Kruschwitz, A. Caputo (Eds.), Advances in Information Retrieval - 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2-6, 2023, Proceedings, Part I, volume 13980 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 33–48. URL: https://doi.org/10.1007/978-3-031-28244-7_3. doi:10.1007/978-3-031-28244-7_3.

[16] L. Sang, M. Xu, S. Qian, X. Wu, Knowledge graph enhanced neural collaborative recommendation, Expert Syst. Appl. 164 (2021) 113992. URL: https://doi.org/10.1016/j.eswa.2020.113992. doi:10.1016/j.eswa.2020.113992.

[17] T. Wang, D. Shi, Z. Wang, S. Xu, H. Xu, Mrp2rec: Exploring multiple-step relation path semantics for

knowledge graph-based recommendations, IEEE Access 8 (2020) 134817–134825. URL: https://doi.org/10.1109/ACCESS.2020.3011279. doi:10.1109/ACCESS.2020.3011279.

[18] D. Shi, T. Wang, H. Xing, H. Xu, A learning path recommendation model based on a multidimensional knowledge graph framework for e-learning, Knowl. Based Syst. 195 (2020) 105618. URL: https://doi.org/10.1016/j.knosys.2020.105618. doi:10.1016/j.knosys.2020.105618.

[19] A. C. M. Mancino, A. Ferrara, S. Bufi, D. Malitesta, T. D. Noia, E. D. Sciascio, Kgtore: Tailored recommendations through knowledge-aware gnn models, in: RecSys '23: Seventeenth ACM Conference on Recommender Systems, Singapore, Singapore, 18 September 2023 - 22 September 2023, ACM, 2023. URL: https://doi.org/10.1145/3604915.3608804. doi:10.1145/3604915.3608804.

[20] V. Paparella, V. W. Anelli, L. Boratto, T. D. Noia, Reproducibility of multi-objective reinforcement learning recommendation: Interplay between effectiveness and beyond-accuracy perspectives, in: RecSys '23: Seventeenth ACM Conference on Recommender Systems, Singapore, Singapore, 18 September 2023 - 22 September 2023, ACM, 2023. URL: https://doi.org/10.1145/3604915.3609493. doi:10.1145/3604915.3609493.

[21] V. Paparella, V. W. Anelli, F. M. Nardini, R. Perego, T. D. Noia, Post-hoc selection of pareto-optimal solutions in search and recommendation, in: CIKM '23: 32nd ACM International Conference on Information and Knowledge Management, Birmingham, United Kingdom, 21 October 2023 - 25 October 2023, ACM, 2023. URL: https://doi.org/10.1145/3583780.3615010. doi:10.1145/3583780.3615010.

[22] D. Vrandecic, M. Krötzsch, Wikidata: a free collaborative knowledgebase, Commun. ACM 57 (2014) 78–85. URL: https://doi.org/10.1145/2629489. doi:10.1145/2629489.

[23] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, C. Bizer, Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia, Semantic Web 6 (2015) 167–195. URL: https://doi.org/10.3233/SW-140134. doi:10.3233/SW-140134.

[24] K. D. Bollacker, C. Evans, P. Paritosh, T. Sturge, J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: J. T. Wang (Ed.), Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008, ACM, 2008, pp. 1247–1250. URL: https://doi.org/10.1145/1376616.1376746. doi:10.1145/1376616.1376746.

[25] Z. Sun, D. Yu, H. Fang, J. Yang, X. Qu, J. Zhang, C. Geng, Are we evaluating rigorously? benchmarking recommendation for reproducible evaluation and fair comparison, in: R. L. T. Santos, L. B. Marinho, E. M. Daly, L. Chen, K. Falk, N. Koenigstein, E. S. de Moura (Eds.), RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020, ACM, 2020, pp. 23–32. URL: https://doi.org/10.1145/3383313.3412489. doi:10.1145/3383313.3412489.

[26] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, TiiS 5 (2016) 19:1–19:19.

[27] M. Richardson, R. Agrawal, P. M. Domingos, Trust management for the semantic web, in: D. Fensel, K. P. Sycara, J. Mylopoulos (Eds.), The Semantic Web - ISWC 2003, Second International Semantic Web Conference, Sanibel Island, FL, USA, October 20-23, 2003, Proceedings, volume 2870 of *Lecture Notes in Computer Science*, Springer, 2003, pp. 351–368. URL: https://doi.org/10.1007/978-3-540-39718-2_23. doi:10.1007/978-3-540-39718-2_23.

[28] C. Ziegler, S. M. McNee, J. A. Konstan, G. Lausen, Improving recommendation lists through topic diversification, in: A. Ellis, T. Hagino (Eds.), Proceedings of the 14th international conference on World Wide Web, WWW 2005, Chiba, Japan, May 10-14, 2005, ACM, 2005, pp. 22–32. URL: https://doi.org/10.1145/1060745.1060754. doi:10.1145/1060745.1060754.

[29] G. Piao, J. G. Breslin, Transfer learning for item recommendations and knowledge graph completion in item related domains via a co-factorization model, in: A. Gangemi, R. Navigli, M. Vidal, P. Hitzler, R. Troncy, L. Hollink, A. Tordai, M. Alam (Eds.), The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings, volume 10843 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 496–511. URL: https://doi.org/10.1007/978-3-319-93417-4_32. doi:10.1007/978-3-319-93417-4_32.

[30] I. Vagliano, D. Monti, A. Scherp, M. Morisio, Content recommendation through semantic annotation of user reviews and linked data, in: Ó. Corcho, K. Janowicz, G. Rizzo, I. Tiddi, D. Garijo (Eds.), Proceedings of the Knowledge Capture Conference, K-CAP 2017, Austin, TX, USA, December 4-6, 2017, ACM, 2017, pp. 32:1–32:4. URL: https://doi.org/10.1145/3148011.3148035. doi:10.1145/3148011.3148035.

[31] M. Alshammari, O. Nasraoui, Augmented semantic explanations for collaborative filtering recommendations, in: A. L. N. Fred, J. Filipe (Eds.), Proceedings of the 11th International Joint Conference on Knowledge Discovery, Knowledge Engi-

neering and Knowledge Management, IC3K 2019, Volume 1: KDIR, Vienna, Austria, September 17-19, 2019, ScitePress, 2019, pp. 83–88. URL: https://doi.org/10.5220/0008070900830088. doi:10.5220/0008070900830088.

[32] W. X. Zhao, G. He, K. Yang, H. Dou, J. Huang, S. Ouyang, J. Wen, Kb4rec: A data set for linking knowledge bases with recommender systems, Data Intell. 1 (2019) 121–136. URL: https://doi.org/10.1162/dint_a_00008. doi:10.1162/dint_a_00008.

[33] X. Wang, T. Huang, D. Wang, Y. Yuan, Z. Liu, X. He, T. Chua, Learning intents behind interactions with knowledge graph for recommendation, in: J. Leskovec, M. Grobelnik, M. Najork, J. Tang, L. Zia (Eds.), WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021, ACM / IW3C2, 2021, pp. 878–887. URL: https://doi.org/10.1145/3442381.3450133. doi:10.1145/3442381.3450133.

[34] W. Chen, P. Huang, J. Xu, X. Guo, C. Guo, F. Sun, C. Li, A. Pfadler, H. Zhao, B. Zhao, POG: personalized outfit generation for fashion recommendation at alibaba ifashion, in: A. Teredesai, V. Kumar, Y. Li, R. Rosales, E. Terzi, G. Karypis (Eds.), Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019, ACM, 2019, pp. 2662–2670. URL: https://doi.org/10.1145/3292500.3330652. doi:10.1145/3292500.3330652.

[35] Z. Sun, Z. Deng, J. Nie, J. Tang, Rotate: Knowledge graph embedding by relational rotation in complex space, in: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019, OpenReview.net, 2019. URL: https://openreview.net/forum?id=HkgEQnRqYQ.

[36] A. Bordes, N. Usunier, A. García-Durán, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: C. J. C. Burges, L. Bottou, Z. Ghahramani, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States, 2013, pp. 2787–2795. URL: https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html.

[37] F. Zhang, N. J. Yuan, D. Lian, X. Xie, W. Ma, Collaborative knowledge base embedding for recommender systems, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016, pp. 353–362. URL: https://doi.org/10.1145/2939672.2939673. doi:10.1145/2939672.2939673.

[38] A. Uyar, F. M. Aliyu, Evaluating search features of google knowledge graph and bing satori: Entity types, list searches and query interfaces, Online Inf. Rev. 39 (2015) 197–213. URL: https://doi.org/10.1108/OIR-10-2014-0257. doi:10.1108/OIR-10-2014-0257.

[39] X. Wang, D. Wang, C. Xu, X. He, Y. Cao, T. Chua, Explainable reasoning over knowledge graphs for recommendation, in: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019, AAAI Press, 2019, pp. 5329–5336. URL: https://doi.org/10.1609/aaai.v33i01.33015329. doi:10.1609/aaai.v33i01.33015329.

[40] V. C. Ostuni, T. D. Noia, E. D. Sciascio, R. Mirizzi, Top-n recommendations from implicit feedback leveraging linked open data, in: Q. Yang, I. King, Q. Li, P. Pu, G. Karypis (Eds.), Seventh ACM Conference on Recommender Systems, RecSys '13, Hong Kong, China, October 12-16, 2013, ACM, 2013, pp. 85–92. URL: https://doi.org/10.1145/2507157.2507172. doi:10.1145/2507157.2507172.

[41] W. X. Zhao, G. He, H. Dou, J. Huang, S. Ouyang, J. Wen, Kb4rec: A dataset for linking knowledge bases with recommender systems, in: A. Cuzzocrea, F. Bonchi, D. Gunopulos (Eds.), Proceedings of the CIKM 2018 Workshops co-located with 27th ACM International Conference on Information and Knowledge Management (CIKM 2018), Torino, Italy, October 22, 2018, volume 2482 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018. URL: http://ceur-ws.org/Vol-2482/paper25.pdf.

[42] T. Di Noia, V. C. Ostuni, P. Tomeo, E. Di Sciascio, Sprank: Semantic path-based ranking for top-N recommendations using linked open data, ACM TIST 8 (2016) 9:1–9:34. URL: https://doi.org/10.1145/2899005. doi:10.1145/2899005.

[43] A. Ragone, P. Tomeo, C. Magarelli, T. D. Noia, M. Palmonari, A. Maurino, E. D. Sciascio, Schema-summarization in linked-data-based feature selection for recommender systems, in: A. Seffah, B. Penzenstadler, C. Alves, X. Peng (Eds.), Proceedings of the Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, April 3-7, 2017, ACM, 2017, pp. 330–335. URL: https://doi.org/10.1145/3019612.3019837. doi:10.1145/3019612.3019837.

[44] V. W. Anelli, T. D. Noia, E. D. Sciascio, A. Fer-

rara, A. C. M. Mancino, Sparse feature factorization for recommender systems with knowledge graphs, in: H. J. C. Pampín, M. A. Larson, M. C. Willemsen, J. A. Konstan, J. J. McAuley, J. Garcia-Gathright, B. Huurnink, E. Oldridge (Eds.), RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September 2021 - 1 October 2021, ACM, 2021, pp. 154–165. URL: https://doi.org/10.1145/3460231.3474243. doi:10.1145/3460231.3474243.

[45] T. Di Noia, Knowledge-enabled recommender systems: Models, challenges, solutions, in: KDWeb, volume 1959 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2017.

[46] M. A. Jaro, Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida, Journal of the American Statistical Association 84 (1989) 414–420. URL: https://www.tandfonline.com/doi/abs/10.1080/01621459.1989.10478785. doi:10.1080/01621459.1989.10478785. arXiv:https://www.tandfonline.com/doi/pdf/10.1080/01621459.1989.10478785.

[47] R. A. Wagner, M. J. Fischer, The string-to-string correction problem, J. ACM 21 (1974) 168–173.

[48] G. Kondrak, *N*-gram similarity and distance, in: SPIRE, volume 3772 of *Lecture Notes in Computer Science*, Springer, 2005, pp. 115–126.

[49] A. Moro, A. Raganato, R. Navigli, Entity linking meets word sense disambiguation: a unified approach, Trans. Assoc. Comput. Linguistics 2 (2014) 231–244. URL: https://tacl2013.cs.columbia.edu/ojs/index.php/tacl/article/view/291.

[50] P. N. Mendes, M. Jakob, A. García-Silva, C. Bizer, Dbpedia spotlight: shedding light on the web of documents, in: C. Ghidini, A. N. Ngomo, S. N. Lindstaedt, T. Pellegrini (Eds.), Proceedings the 7th International Conference on Semantic Systems, I-SEMANTICS 2011, Graz, Austria, September 7-9, 2011, ACM International Conference Proceeding Series, ACM, 2011, pp. 1–8. URL: https://doi.org/10.1145/2063518.2063519. doi:10.1145/2063518.2063519.

[51] G. Rizzo, R. Troncy, NERD: A framework for unifying named entity recognition and disambiguation extraction tools, in: W. Daelemans, M. Lapata, L. Màrquez (Eds.), EACL 2012, 13th Conference of the European Chapter of the Association for Computational Linguistics, Avignon, France, April 23-27, 2012, The Association for Computer Linguistics, 2012, pp. 73–76. URL: https://aclanthology.org/E12-2015/.

[52] J. Daiber, M. Jakob, C. Hokamp, P. N. Mendes, Improving efficiency and accuracy in multilingual entity extraction, in: M. Sabou, E. Blomqvist, T. D. Noia, H. Sack, T. Pellegrini (Eds.), I-SEMANTICS 2013 - 9th International Conference on Semantic Systems, ISEM '13, Graz, Austria, September 4-6, 2013, ACM, 2013, pp. 121–124. URL: https://doi.org/10.1145/2506182.2506198. doi:10.1145/2506182.2506198.