The Yelp Collaborative Knowledge Graph

Theis E. Jendal
tjendal@cs.aau.dk
Aalborg University
Denmark

Mads Corfixen
maco@cs.aau.dk
Aalborg University
Denmark

Magnus Olesen
magnus_olesen@live.dk
Aalborg University
Denmark

Peter Dolog
dolog@cs.aau.dk
Aalborg University
Denmark

Katja Hose
katja.hose@tuwien.ac.at
TU Wien
Austria

Daniele Dell'Aglio
dade@cs.aau.dk
Aalborg University
Denmark

Matteo Lissandrini
matteo.lissandrini@univr.it
Verona University
Italy

Abstract

Yelp Open Dataset (YOD) is a widely used dataset for Recommender Systems (RS). Multiple Knowledge Graphs (KGs) have been built for YOD, but they have various issues: the conversion processes usually do not follow state-of-the-art methodologies, fail to properly link to other KGs, do not link to existing vocabularies, ignore important data, and are generally of small size. Instead, we present the Yelp Collaborative Knowledge Graph (YCKG), where we correctly integrating taxonomies, product categories, business locations, and the Yelp social network, through common practices within the semantic web community, overcoming all these issues. As a result, the YCKG includes 150k businesses and 16.9M reviews from 1.9M distinct real users, resulting in over 244 million triples, 144 distinct predicates, for about 72 million resources, with an average in-degree and out-degree of 3.3 and 12.2, respectively. Further, we release both the data and the code used to generate the KG for inspection and further extensions. This dataset can be used to develop and test both recommendation and data-mining algorithms able to exploit rich and semantically meaningful knowledge. We publicize the code for the CKG construction on: https://github.com/MadsCorfixen/The-Yelp-Collaborative-Knowledge-Graph.

CCS Concepts

• **Information systems → Test collections**; **Relational database model**; *Open source software*; **Collaborative filtering**; **Social recommendation**; **Ontologies**.

Keywords

Dataset, Recommender systems, Knowledge graph, Open data

ACM Reference Format:
Theis E. Jendal, Mads Corfixen, Magnus Olesen, Peter Dolog, Katja Hose, Daniele Dell'Aglio, and Matteo Lissandrini. 2025. The Yelp Collaborative Knowledge Graph. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25), November 10–14, 2025, Seoul, Republic of Korea.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3746252.3761615

ACM ISBN 979-8-4007-2040-6/2025/11
https://doi.org/10.1145/3746252.3761615

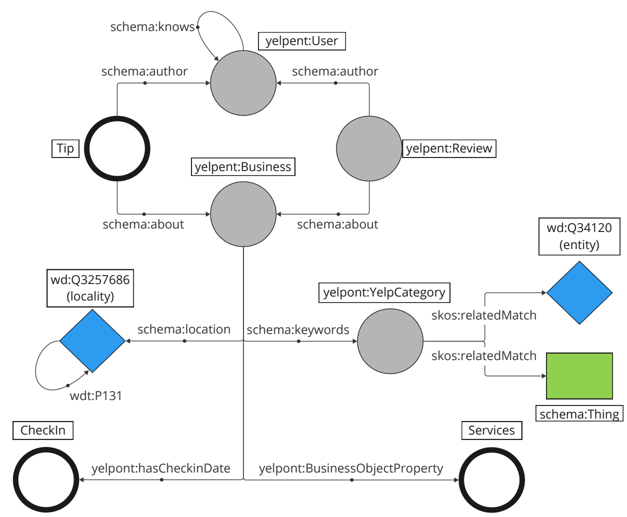

Figure 1: The structure of the YCKG. Circles represent Yelp IRIs, non-filled circles blank nodes, squares Schema.org IRIs, and diamonds Wikidata IRIs.

1 Introduction

In the Recommender System (RS) domain, there is an abundance of datasets describing products, businesses, and their customer relationships (e.g., ratings and reviews). Such data is exploited to build better services powered by intelligent recommendations. Many RSs adopt Collaborate Filtering (CF) [3, 10, 15, 41, 42], which assumes users with similar past behavior have similar preferences; more advanced methods extend this by using contextual information, such as taxonomies or item descriptions to handle sparse interactions [2, 12, 21, 22, 41]. With the increasing popularity of KGs, the RS community gained interest in exploiting KGs as a technology to organize contextual information [2, 12]. In these cases, the collaborative information is modeled as a bipartite graph with users, products, and their interactions. Moreover, the products are enriched with connections between themselves, and to non-recommendable entities (e.g., categories or descriptive entities), resulting in a Collaborative Knowledge Graph (CKG). However, none of the existing datasets fully utilize standard Semantic Web methodologies when creating a domain-specific KG to test RS (Table 1), with no standard ontologies or vocabularies being adopted or defined. This has been the case also for the Yelp Open Dataset (YOD) [46]. In the

Table 1: Existing CKGs: showing the statistics of the most sound. We use '*' to represent availability of the original collaborative information and CKG. The method used for the CKG construction: crawling N-Hops connections from aligned entities, Hand Selected (HS) relations, or converting Attributes to nodes instead linking to an External KG. #U, #I, #E, and #P are the number of users, items, entities (w/o item and review nodes), and predicates, respectively. We define three levels of soundness: ✔utilizes Semantic Web approaches for construction; (✔) connected to an open-domain KG; and ✗no theory or open-domain KG adoption;

Domain	Original data	CKG	Ext. KG	Method	#U	#I	#E	#P	Sound
Movie	MovieLens20m* [13]	[17]	Freebase	1-Hop	61,583	19,533	1,125,100	81	(✔)
		[38]*	Satori	1-Hop	138,159	16,954	102,569	32	(✔)
		[18]*	Freebase	1-Hop	138,287	13,047	40,529	66	(✔)
	MindReader* [3]	[3]*	WikiData	HS	1,174	9,000	10,030	8	(✔)
Book	Book-Crossing* [51]	[38]	Satori	1-Hop	70,679	24,915	88,572	39	(✔)
		[36]*	Satori	4-Hop	17,860	14,967	–	–	(✔)
	Amazon Book* [14]	[50]*	Freebase	2-Hop	70,679	24,915	88,572	39	(✔)
	DBbook2014* [27]	[4]*	DBPedia	1-Hop	5,576	2,680	13,882	13	(✔)
Music	Last.FM* [31]	[50]*	Freebase	2-Hop	23,566	48,123	58,266	9	(✔)
		[38]*	Satori	1-Hop	1,872	3,846	9,366	60	(✔)
POI	Yelp2018* [46]	[45]*	Freebase	2-hop	45,919	45,538	47,472	42	(✔)
		[41]*	–	Attribute	45,919	45,538	90,961	42	
		[48]	–	Attribute	43,873	11,537	285,317	4	
		[49]*	–	Attribute	36,105	22,496	1121	9	✗
		[35]	–	Attribute	37,940	11,516	46,606	7	
		[16]	–	Attribute	16,239	14,284	558	4	
		Ours*	**Schema & Wikidata**	**HS**	**1,987,897**	**150,346**	**67,238,177**	**267**	✔

Table 2: Yelp Open Dataset JSON files.

File	Size	#Objects
Business	113 MB	150,346
Checkin	274 MB	131,930
Review	4,980 MB	6,990,280
Tip	172 MB	908,915
User	3,130 MB	1,987,897

Figure 2: Popular business over time.

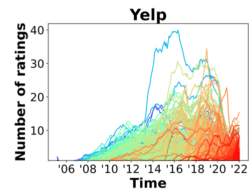

past, the Yelp2018 dataset [35, 41, 48] has already been converted to a CKG, to evaluate recommender systems that can exploit data represented as a graph. Unfortunately, as we describe below, past approaches, the original JSON dump has been directly converted to a set of triples, meaning the JSON node represents the whole JSON object instead of its constituents. It does not preserve all the original data, ignoring e.g. tips and social information. Among others, the heterogeneous graphs constructed with attribute-based methods omit a lot of semantic information. Existing KGs constructed by methods based on k-hops, on the other hand, have poor coverage of entities of the original Yelp dataset. To overcome these limitations, we release Y2KG [9], a tool to create the Yelp Collaborative KG (YCKG), a KG for the YOD, which contains data on more than 8 million entities, mostly businesses and reviews, while preserving almost all the metadata and semantic descriptions attached to them (see Figure 1). Thus, we present the first conversion of the YOD to a domain-specific KG, providing full coverage of users and items, the YCKG, while aiming to adhere to standards and best practices.

In this work, we thus present: (i) a sound methodology and a tool for converting the YOD to a CKG, (ii) an ontology for the YOD, and (iii) a methodology for mapping categories and locations to entities. The resulting open dataset [8], thus, features: (a) the first CKG released in standard-compliant RDF based on semantic design principles, (b) the first CKG with social information, and (c) the largest CKG for YOD. Specifically, our graph has over 244 million triples, approximately 72 million resources, 144 predicates, with the average in-degree and out-degrees being 3.35 and 12.20 respectively. Linking 310 Yelp categories to *schema.org* types, 136 of which linked to Wikidata entities, and further linking businesses to entities representing cities, counties, and states.

2 Related Work

KGs combined with user-item interactions, referred to as CKGs, provide contextual information connecting rated items to other descriptive entities [3, 32, 35, 36, 41]. A single interaction, therefore, provides information not only about the individual item but also the relations to other similar items and thus it is possible to derive also information about users in sparse settings [41]. Some methods exploiting KGs are designed to provide explanations for increased user trust in the recommendation, such as "...because the user likes Italian restaurants with private parking", which can be described as paths and relations in a KG [18, 34, 41].

Yet, most existing recommendation benchmark datasets do not have a KG, e.g., Amazon Book [14], DBBook2014 [27], Yelp [46], Last FM [31], and Movielens [13]. Hence, approaches using KG information often construct their own KG [7, 12]. The KG creation processes utilized rarely adopt state-of-the-art methodologies and design principles. Table 1 is a summary of related KG datasets and the representative methods utilizing those datasets. Detailed overviews can be found in related works [7, 12]. We find that none of the existing datasets fully adopt standard Semantic Web methodologies when creating a domain-specific KG, e.g., no standard ontology or vocabulary is adopted or defined. We highlight that the version of KGs for the Yelp Open Dataset (YOD) currently used by most methods suffers from the following issues: (i) containing few triples (approximately 2 million[41]) while YOD contains almost 7 million reviews about more than 130K businesses; (ii) disregard social network information; (iii) does not link to any open-domain KG[1]; and (iv) treat serialized JSON objects as single node identifiers, disregarding in this way the rich set of attributes they contain. This results in a large information loss and to widely varying numbers of users, items, entities, and relations across datasets, making subsequent evaluations across methods difficult or impossible.

Therefore without a high-quality KGs it is practically impossible to study and test advanced methods able to exploit such rich data. Given the YOD popularity in the RS community [15, 25, 42] and its frequent adoption for evaluating RS with KG information [6, 32, 33,

[1]Freebase has remained unmaintained since 2015 [11].

Table 3: Validation of the competency questions.

	Competency Question	Answer	Validated?
1	How many different types of businesses are defined in Yelp?	1,311	Yes
2	How many businesses of type "Restaurants" exist?	52,268	Yes
3	How many businesses of type "Restaurants" have been reviewed?	52,268	Yes
4	How many businesses have been reviewed?	150,346	Yes
5	How many businesses have, on average, a rating above 4?	43,499	Yes
6	What is the average rating across businesses?	3.5957	Yes
7	How many businesses have been reviewed in Santa Barbara, CA?	3,829	Yes
8	What business has the highest number of visits in one day?	465	Yes
9	How many visits does the most visited business have?	52,144	Yes
10	How many people have written a review on Yelp?	1,987,929	Yes
11	How many users have more than 10 friends?	758,803	Yes
12	How many friends does a user have on average?	52.93	Yes
13	How many users have authored 10 reviews?	14,119	Yes
14	How many reviews did users make in May 2018?	79,434	Yes
15	How many parking options can a business provide?	5	Yes
16	How many businesses have karaoke?	75	Yes
17	How many people live in Indianapolis?	887,642	Yes
18	How many cities are in the state of Florida?	391	Yes

35, 41, 48, 49], we select it as a focus of this work. Different from existing YOD-based KGs, we establish clear competency questions to define the requirements for the KG [1] and align the KG with open-domain ontologies, taxonomies, and entities. Obtaining the largest YOD-based KG for testing KG-powered RS.

3 Constructing the YCKG

The YOD dump contains multiple files describing Businesses, Users, and Reviews (Table 2), each having different characteristics. To capture the characteristics in our constructed CKG, we define competency questions, the final schema, and the structure of the YCKG.

Competency Questions. To guide our design of the desired target schema in which to convert the YOD, we listed a series of 16 competency questions [1]. These competency questions have been designed to ensure that no data is lost during the data conversion process while also guiding the integration of external data. They are supposed to be very simple questions to guide the data modeling, while our resulting resource can both express and support answers to more complex and advanced data mining tasks and research applications. We then used these competency questions to verify the appropriateness and quality of the CKG and of the data conversion process. This provides us the guarantee that all the relevant data has been converted successfully and that we can identify contextual information about businesses and user social networks and preferences. We define the questions in Table 3 and validate that our CKG is capable of answering them.

Conversion to RDF. Given the requirements derived from the competency questions above, we design the Y2KG [9], a tool to losslessly convert the YOD, which is a set of five newline-delimited JSON files, into a KG serialized as RDF triples in a set of N-Triple files. The structure of the YCKG is shown in Figure 1.

The JSON files contain a distinct object on each line, so that each entity of interest corresponds to a line in the file with the first key/value pair indicating the object ID. This ID is used to create an IRI for the User, Business, and Review nodes. As neither Tips nor Checkins have IDs, we represent them as blank nodes. Each Yelp Entity Business, Review, and Checkin node is then associated, through `rdf:type`, to Schema.org classes' `LocalBusiness`, `UserReview`, and `ArriveAction`. As Schema.org does not include a class for Tip, we define it in the Yelp Vocabulary [47].

The keys in each key/value pair in the JSON objects are usually mapped to RDF properties, and their values to RDF terms, which can be either literals or resources. We have manually inspected all the properties and associated them to Schema.org properties whenever possible [2]. When the attribute values are strings representing lists or maps, we map them to blank nodes annotated by properties. In the dataset, there are a total of 8 different distinct types of attributes that in the conversion process are mapped to as many distinct types of blank nodes. For categories, we instantiate a URI for every category in the Yelp dataset, for a total of 1, 304 distinct business categories. For attributes for which we could not find a match in *Schema.org*, we create a property in the Yelp vocabulary instead. In total we have re-used 22 properties from the Schema.org vocabulary, and 116 are created in the Yelp vocabulary.

4 Enriching the YCKG

The YOD provides mentions to categories and products that can be mapped to Wikidata entities or Schema.org concepts. The following describes the automatic process used for this mapping.

Mapping Yelp Categories to Schema.org. We identify candidate matches between Yelp categories and Schema.org terms by using word embeddings. In practice, we use a sentence transformer [19, 30] to map both labels of Yelp categories and labels of the Schema.org types to the same 384-dimensional vector space. Then, given a label of a Yelp category and the corresponding vector, the *Schema.org* type with the highest cosine similarity of its vectors is chosen as the candidate mapping.

Mappings to Wikidata. We also link the YCKG to Wikidata entities, both for the case of business categories and for business locations, allowing users of YCKG to use LLMs by querying Wikipedia for textual descriptions [28]. To map business categories, we exploit the information already present in Wikidata. Specifically, Wikidata includes some mappings to Schema.org resources using the "*equivalent class*" property (`wdt:P1709`). We further link the Yelp locations to Wikidata locations. Hence, we start by collecting from the YOD all the distinct (city, state)-pairs and the associated geo-location of the businesses. We group businesses by the (city, state)-pair, computing the centroid of the group to get a general geographical location. Then, for each (city, state)-pair, we find potential Wikidata QID matches by submitting a search query to the MediaWiki Wikibase API [26]. Allowing for automatic disambiguation of the pairs in the YOD to Wikidata-locations based on their coordinates.

5 Evaluation and Analysis

We present some relevant statistics extracted from YCKG (Table 4). The YCKG contains more than 242M facts about 10.4M subjects using 144 unique predicates. Around half of the triples are associated to the social network component (through the `schema:knows` relation in Figure 1), which is a rich resource that can be exploited when trying to infer user preferences from the past preferences of their social circle. Related to this, the average outdegree of nodes representing users is 74, since, as shown below (Table 3), we see that on average users have 52 friends; the remaining edges link to

[2] Note that some attributes are nested within a JSON object for the key `attributes`. Since the key itself does not carry any meaningful context for the business, the key/-value pairs are extracted and treated as the other attributes.

Table 4: Basic statistics of the final Yelp knowledge graph.

Statistic	Result
Triples	242,247,823
Unique Subjects	10,495,829
Unique Objects	61,450,383
Unique Predicates	267
Freq. Most Prevalent Predicate (schema:knows)	105,225,474
Freq. Most Prevalent Class (schema:UserReview)	6,990,280
Average Indegree (Overall)	3.35
Average Outdegree (Overall)	12.20
Average Indegree (Businesses)	46.49
Average Outdegree (Businesses)	114.02
Average Indegree (Users)	3.97
Average Outdegree (Users)	74.10

user metadata. Further, we see reviews represent a core piece of the data, since ∼69% of subjects (6.9M) are reviews. Thus, the dataset is extremely well-fitted to train RSs and maintain far more of the original dataset than any existing CKG for YOD (Table 1). Hence, Y2KG (i) preserves orders of magnitude more data than previous efforts, and (ii) introduces high-quality mappings to external resources.

Soundness of the Conversion Script. We use the competency questions (Table 3), released as part of the source code [9], to validate the soundness of the YCKG. The correct results are based on the original dataset and compared to the corresponding SPARQL queries on our CKG. These results illustrate the higher quality of the YCKG compared to the other CKGs outlined in section 2, as these do not preserve the data to answer the competency questions. Interestingly, the results also provide some insights into the contents of the dataset. For instance: there are more than 1.3k categories on Yelp, with ∼30% of businesses being restaurants); the average rating is approximately ∼3.6 while less than 29% of businesses have a rating above 4.0; and the social network portion of the graph is also well connected, with users having on average ∼53 friends and around 38% of users having more than 10 friends. The YCKG is also one of the few existing datasets that provides timestamps to user-item interactions. In Figure 2, we analyze the number of ratings for the most popular businesses. We see that the number of reviews increases over time, but that the data is also affected by the COVID-19 pandemic. This pattern suggests that the dataset reflects real-world dynamics, with organically increasing user engagement interrupted by the pandemic.

Evaluating Schema.org Mappings. To ensure the quality of mapped entities, we compute a minimum threshold below which candidate matches are discarded. To estimate the threshold, we sampled and manually mapped 200 random pairs for ground truth labels. By analyzing the recall, precision, and F_1 score, we observed that a similarity above 0.91 are sufficiently reliable to be accepted automatically, which corresponds to 94 mappings to be included in this way. We discarded mappings for 927 Yelp categories and manually verified the remaining 291. Overall, this results in 23.6% of all Yelp categories being mapped to a Schema.org term. Because successfully mapped categories are also the most frequent, such as the Restaurant category, our mappings ensure that 94.6% of businesses (items) have at least one category mapped to a Schema.org term and 85.7% to at least one category mapped to a Wikidata concept.

Evaluating Wikidata Location Mappings. We sample some business locations to estimate the quality of this mapping. In our sample, we achieve precision at the city level of 69%. Our analysis of the incorrect mappings is grounded in incorrect data in the original YOD, often due to misspellings. Moreover, in our sample we obtain a recall of about 70% since 30% of the corresponding cities are not present in Wikidata. Nonetheless, we have found that across the 1467 unique $(city, state)$-pairs, 783 of them appear only once (i.e., associated with only one business) in the entire dataset. Meaning that the precision does not actually indicate a widespread issue. Focusing only on $(city, state)$-pairs that appear more than once in the dataset, then 99.71% of businesses are assigned to a correct location. Therefore, this dataset present high quality temporal and spatial information for the first time mapped to a high-quality KG.

Table 5: Related Yelp CKG statistics.

Dataset	#Reviews	#Relations	CG Density	CKG Density
[45]	1,183,610	869,603	1.42e-04	1.06e-04
[41]	1,185,068	1,853,704	1.42e-04	1.62e-04
[48]	229,907	570,634	7.49e-05	4.44e-06
[49]	191,506	467,288	5.58e-05	7.40e-06
[35]	229,178	302,937	9.37e-05	5.77e-05
[16]	198,397	54,276	2.13e-04	2.62e-04
Ours	**6,990,280**	**242,247,823**	1.53e-06	5.36e-08

Use-cases. The dataset has been used in the benchmarking of recommendation approaches [22, 23]. Here it is shown that even purely collaborative-based methods struggle when working on the larger dataset. For example, the best performing model INMO [43] has NDCG@20 0.035 for YCKG [22], while on a smaller Yelp dataset it is 0.065 [43]. This illustrates that more work is needed to make use of the extended dataset in recommender systems. In Table 5, we show that YCKG is far larger than the related works' KG. While the proposed use-case for this dataset is to study recommender systems, it is not limited to this domain. Due to the amount of predicates, node types, and locations, it can be used for node-classification [44], link-prediction [39], POI analysis [20], and LLM improvements [28].

6 Conclusion and Future Work

Existing approaches disregard large portions of the Yelp Open Dataset [46] when converting it to a KG, and thus do not present realistic results and miss important opportunities. For this reason, we present here a tool to convert the YOD from a JSON dump to a standards-compliant knowledge graph in RDF, as well as an ontology for describing the Yelp Open Dataset. Our process allows for the first time the full conversion of the abundant information available in the original dataset. Moreover, it allows for the first time to exploit external knowledge from Schema.org (and its derived taxonomy) and Wikidata to enrich the original data dump with contextual knowledge. In the future, we aim to (a) study methods to extend the quality and coverage of our mappings, (b) extend the adopted technologies like SHACL for validation and `rml.io` for more declarative conversion, and (c) adopt this dataset for a re-evaluation campaign of state-of-the-art recommendation methods.

Acknowledgments

This research was partially funded by the EU's H2020 research innovation programme under grant agreement No 838216, the Poul Due Jensen Foundation, and the Independent Research Fund Denmark (DFF) under grant agreement no. DFF-8048-00051B. We are thankful to Christian F.P. Nielsen and Thomas Heede for their substantial contribution to the initial version of the Y2KG tool.

GenAI Usage Disclosure

Co-pilot was used during code writing. Generative AI has been used for spell checking and grammar, with the initial draft being written without the use of any tools.

References

[1] Camila Bezerra, Fred Freitas, and Filipe Santana. 2013. Evaluating ontologies with competency questions. In *2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, Vol. 3. IEEE, 284–285.

[2] Russa Biswas, Lucie-Aimée Kaffee, Michael Cochez, Stefania Dumbrava, Theis E. Jendal, Matteo Lissandrini, Vanessa López, Eneldo Loza Mencía, Heiko Paulheim, Harald Sack, Edlira Vakaj, and Gerard de Melo. 2023. Knowledge Graph Embeddings: Open Challenges and Opportunities. *TGDK* 1, 1 (2023), 4:1–4:32. doi:10.4230/TGDK.1.1.4

[3] Anders H Brams, Anders L Jakobsen, Theis E Jendal, Matteo Lissandrini, Peter Dolog, and Katja Hose. 2020. MindReader: recommendation over knowledge graph entities with explicit user ratings. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2975–2982.

[4] Yixin Cao, Xiang Wang, Xiangnan He, Zikun Hu, and Tat-Seng Chua. 2019. Unifying Knowledge Graph Learning and Recommendation: Towards a Better Understanding of User Preferences. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 151–161. doi:10.1145/3308558.3313705

[5] Tommaso Carraro, Alessandro Daniele, Fabio Aiolli, and Luciano Serafini. 2022. Logic Tensor Networks for Top-N Recommendation. In *AIxIA 2022 - Advances in Artificial Intelligence - XXIst International Conference of the Italian Association for Artificial Intelligence, AIxIA 2022, Udine, Italy, November 28 - December 2, 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13796)*, Agostino Dovier, Angelo Montanari, and Andrea Orlandini (Eds.). Springer, 110–123. doi:10.1007/978-3-031-27181-6_8

[6] Rose Catherine and William W. Cohen. 2016. Personalized Recommendations using Knowledge Graphs: A Probabilistic Logic Programming Approach. In *Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15-19, 2016*, Shilad Sen, Werner Geyer, Jill Freyne, and Pablo Castells (Eds.). ACM, 325–332. doi:10.1145/2959100.2959131

[7] Janneth Chicaiza and Priscila Valdiviezo Díaz. 2021. A Comprehensive Survey of Knowledge Graph-Based Recommender Systems: Technologies, Development, and Contributions. *Inf.* 12, 6 (2021), 232. doi:10.3390/info12060232

[8] Mads Corfixen, Magnus Olesen, Thomas Heede, and Christian Filip Pinderup Nielsen. 2023. The Yelp Collaborative Knowledge Graph. doi:10.5281/zenodo.7878446

[9] Corfixen, Mads. 2023. *The Yelp Collaborative Knowledge Graph.* https://github.com/MadsCorfixen/The-Yelp-Collaborative-Knowledge-Graph Accessed 27 Apr. 2025.

[10] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of recommender algorithms on top-n recommendation tasks. In *RecSys'10*.

[11] Google. [n. d.]. *Freebase API.* https://developers.google.com/freebase Accessed 13 June 2025.

[12] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2022. A Survey on Knowledge Graph-Based Recommender Systems. *IEEE Trans. Knowl. Data Eng.* 34, 8 (2022), 3549–3568. doi:10.1109/TKDE.2020.3028705

[13] F. Maxwell Harper and Joseph A. Konstan. 2016. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4 (2016), 19:1–19:19. doi:10.1145/2827872

[14] Ruining He and Julian J. McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*, Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao (Eds.). ACM, 507–517. doi:10.1145/2872427.2883037

[15] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR'20*.

[16] Binbin Hu, Chuan Shi, Wayne Xin Zhao, and Philip S. Yu. 2018. Leveraging Meta-path based Context for Top- N Recommendation with A Neural Co-Attention Model. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 1531–1540. doi:10.1145/3219819.3219965

[17] Jin Huang, Wayne Xin Zhao, Hongjian Dou, Ji-Rong Wen, and Edward Y. Chang. 2018. Improving Sequential Recommendation with Knowledge-Enhanced Memory Networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, Kevyn Collins-Thompson, Qiaozhu Mei, Brian D. Davison, Yiqun Liu, and Emine Yilmaz (Eds.). ACM, 505–514. doi:10.1145/3209978.3210017

[18] Xiaowen Huang, Quan Fang, Shengsheng Qian, Jitao Sang, Yan Li, and Changsheng Xu. 2019. Explainable Interaction-driven User Modeling over Knowledge Graph for Sequential Recommendation. In *Proceedings of the 27th ACM International Conference on Multimedia, MM 2019, Nice, France, October 21-25, 2019*, Laurent Amsaleg, Benoit Huet, Martha A. Larson, Guillaume Gravier, Hayley Hung, Chong-Wah Ngo, and Wei Tsang Ooi (Eds.). ACM, 548–556. doi:10.1145/3343031.3350893

[19] HuggingFace. 2023. *all-MiniLM-L6-v2.* https://huggingface.co/optimum/all-MiniLM-L6-v2 Accessed 09 May 2025.

[20] Md. Ashraful Islam, Mir Mahathir Mohammad, Sarkar Snigdha Sarathi Das, and Mohammed Eunus Ali. 2022. A survey on deep learning based Point-of-Interest (POI) recommendations. *Neurocomputing* 472 (2022), 306–325. doi:10.1016/J.NEUCOM.2021.05.114

[21] Theis E. Jendal, Matteo Lissandrini, Peter Dolog, and Katja Hose. 2023. GInRec: A Gated Architecture for Inductive Recommendation using Knowledge Graphs. In *Proceedings of the Fifth Knowledge-aware and Conversational Recommender Systems Workshop co-located with 17th ACM Conference on Recommender Systems (RecSys 2023), Singapore, September 19th, 2023 (CEUR Workshop Proceedings, Vol. 3560)*. CEUR-WS.org, 80–89. https://ceur-ws.org/Vol-3560/long6.pdf

[22] Theis E. Jendal, Matteo Lissandrini, Peter Dolog, and Katja Hose. 2024. The Limits of Graph Samplers for Training Inductive Recommender Systems. *Proc. VLDB Endow.* 18, 8 (2024). doi:10.14778/3742728.3742743

[23] Theis E. Jendal, Matteo Lissandrini, Peter Dolog, and Katja Hose. 2025. Handling new users and items: a comparative study of inductive recommenders. *Data Mining and Knowledge Discovery* 39, 5 (29 Jul 2025), 58. doi:10.1007/s10618-025-01134-2

[24] Federico López, Beatrice Pozzetti, Steve Trettel, Michael Strube, and Anna Wienhard. 2021. Vector-valued Distance and Gyrocalculus on the Space of Symmetric Positive Definite Matrices. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 18350–18366. https://proceedings.neurips.cc/paper/2021/hash/98c39996bf1543e974747a2549b3107c-Abstract.html

[25] Kelong Mao, Jieming Zhu, Jinpeng Wang, Quanyu Dai, Zhenhua Dong, Xi Xiao, and Xiuqiang He. 2021. SimpleX: A simple and strong baseline for collaborative filtering. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1243–1252.

[26] MediaWiki. 2023. *Wikibase API.* https://www.mediawiki.org/wiki/Wikibase/API Accessed 04 May 2025.

[27] Tommaso Di Noia, Iván Cantador, and Vito Claudio Ostuni. 2014. Linked Open Data-Enabled Recommender Systems: ESWC 2014 Challenge on Book Recommendation. In *Semantic Web Evaluation Challenge - SemWebEval 2014 at ESWC 2014, Anissaras, Crete, Greece, May 25-29, 2014, Revised Selected Papers (Communications in Computer and Information Science, Vol. 475)*, Valentina Presutti, Milan Stankovic, Erik Cambria, Iván Cantador, Angelo Di Iorio, Tommaso Di Noia, Christoph Lange, Diego Reforgiato Recupero, and Anna Tordai (Eds.). Springer, 129–143. doi:10.1007/978-3-319-12024-9_17

[28] Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhania, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omeliyanenko, Wen Zhang, Matteo Lissandrini, Russa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, and Damien Graux. 2023. Large Language Models and Knowledge Graphs: Opportunities and Challenges. *TGDK* 1, 1 (2023), 2:1–2:38. doi:10.4230/TGDK.1.1.2

[29] Yanru Qu, Ting Bai, Weinan Zhang, Jian-Yun Nie, and Jian Tang. 2019. An End-to-End Neighborhood-based Interaction Model for Knowledge-enhanced Recommendation. *CoRR* abs/1908.04032 (2019). arXiv:1908.04032 http://arxiv.org/abs/1908.04032

[30] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. https://arxiv.org/abs/1908.10084

[31] Markus Schedl. 2016. The LFM-1b Dataset for Music Retrieval and Recommendation. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval, ICMR 2016, New York, New York, USA, June 6-9, 2016*, John R. Kender, John R. Smith, Jiebo Luo, Susanne Boll, and Winston H. Hsu (Eds.). ACM, 103–110. doi:10.1145/2911996.2912004

[32] Xiao Sha, Zhu Sun, and Jie Zhang. 2021. Hierarchical attentive knowledge graph embedding for personalized recommendation. *Electron. Commer. Res. Appl.* 48 (2021), 101071. doi:10.1016/j.elerap.2021.101071

[33] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and Philip S. Yu. 2019. Heterogeneous Information Network Embedding for Recommendation. *IEEE Trans. Knowl. Data Eng.* 31, 2 (2019), 357–370. doi:10.1109/TKDE.2018.2833443

[34] Weiping Song, Zhijian Duan, Ziqing Yang, Hao Zhu, Ming Zhang, and Jian Tang. 2019. Explainable Knowledge Graph-based Recommendation via Deep Reinforcement Learning. *CoRR* abs/1906.09506 (2019). arXiv:1906.09506 http://arxiv.org/abs/1906.09506

[35] Zhu Sun, Jie Yang, Jie Zhang, Alessandro Bozzon, Long-Kai Huang, and Chi Xu. 2018. Recurrent knowledge graph embedding for effective recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan (Eds.). ACM, 297–305. doi:10.1145/3240323.3240361

[36] Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2018. RippleNet: Propagating User Preferences on the Knowledge Graph for Recommender Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22-26, 2018*, Alfredo Cuzzocrea, James Allan, Norman W. Paton, Divesh Srivastava, Rakesh Agrawal, Andrei Z. Broder, Mohammed J. Zaki, K. Selçuk Candan, Alexandros Labrinidis, Assaf Schuster, and Haixun Wang (Eds.). ACM, 417–426. doi:10.1145/3269206.3271739

[37] Hongwei Wang, Fuzheng Zhang, Mengdi Zhang, Jure Leskovec, Miao Zhao, Wenjie Li, and Zhongyuan Wang. 2019. 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, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 968–977. doi:10.1145/3292500.3330836

[38] Hongwei Wang, Miao Zhao, Xing Xie, Wenjie Li, and Minyi Guo. 2019. Knowledge Graph Convolutional Networks for Recommender Systems. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 3307–3313. doi:10.1145/3308558.3313417

[39] Meihong Wang, Linling Qiu, and Xiaoli Wang. 2021. A Survey on Knowledge Graph Embeddings for Link Prediction. *Symmetry* 13, 3 (2021), 485. doi:10.3390/SYM13030485

[40] Qinqin Wang, Elias Z. Tragos, Neil Hurley, Barry Smyth, Aonghus Lawlor, and Ruihai Dong. 2022. Entity-Enhanced Graph Convolutional Network for Accurate and Explainable Recommendation. In *UMAP '22: 30th ACM Conference on User Modeling, Adaptation and Personalization, Barcelona, Spain, July 4 - 7, 2022*, Alejandro Bellogín, Ludovico Boratto, Olga C. Santos, Liliana Ardissono, and Bart P. Knijnenburg (Eds.). ACM, 79–88. doi:10.1145/3503252.3531316

[41] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *SIGKDD'19*.

[42] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR'19*.

[43] Yunfan Wu, Qi Cao, Huawei Shen, Shuchang Tao, and Xueqi Cheng. 2022. INMO: A Model-Agnostic and Scalable Module for Inductive Collaborative Filtering. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) *(SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 91–101. doi:10.1145/3477495.3532000

[44] Shunxin Xiao, Shiping Wang, Yuanfei Dai, and Wenzhong Guo. 2022. Graph neural networks in node classification: survey and evaluation. *Mach. Vis. Appl.* 33, 1 (2022), 4. doi:10.1007/S00138-021-01251-0

[45] Yuhao Yang, Chao Huang, Lianghao Xia, and Chenliang Li. 2022. Knowledge Graph Contrastive Learning for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) *(SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 1434–1443. doi:10.1145/3477495.3532009

[46] Yelp. 2023. *Yelp Open Dataset*. https://www.yelp.com/dataset Accessed 27 Apr. 2025.

[47] Yelp. 2023. *Yelp Vocabulary*. https://purl.archive.org/purl/yckg/vocabulary Accessed 27 Apr. 2025.

[48] Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. 2014. Personalized entity recommendation: a heterogeneous information network approach. In *Seventh ACM International Conference on Web Search and Data Mining, WSDM 2014, New York, NY, USA, February 24-28, 2014*, Ben Carterette, Fernando Diaz, Carlos Castillo, and Donald Metzler (Eds.). ACM, 283–292. doi:10.1145/2556195.2556259

[49] Huan Zhao, Quanming Yao, Jianda Li, Yangqiu Song, and Dik Lun Lee. 2017. Meta-Graph Based Recommendation Fusion over Heterogeneous Information Networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*. ACM, 635–644. doi:10.1145/3097983.3098063

[50] Wayne Xin Zhao, Gaole He, Kunlin Yang, Hongjian Dou, Jin Huang, Siqi Ouyang, and Ji-Rong Wen. 2019. KB4Rec: A Data Set for Linking Knowledge Bases with Recommender Systems. *Data Intell.* 1, 2 (2019), 121–136. doi:10.1162/dint_a_00008

[51] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web, WWW 2005, Chiba, Japan, May 10-14, 2005*, Allan Ellis and Tatsuya Hagino (Eds.). ACM, 22–32. doi:10.1145/1060745.1060754