

Bridging the User-side Knowledge Gap in Knowledge-aware Recommendations with Large Language Models

Zheng Hu¹, Zhe Li¹, Ziyun Jiao¹, Satoshi Nakagawa²,
Jiawen Deng^{1*}, Shimin Cai¹, Tao Zhou¹, Fuji Ren^{1,3*}

¹School of Computer Science and Engineering, University of Electronic Science and Technology of China

²Graduate School of Information Science and Technology, University of Tokyo

³Shenzhen Institute for Advanced Study, University of Electronic Science and Technology of China
{1998huzheng, dengjw2016, ren2fuji}@gmail.com

Abstract

In recent years, knowledge graphs have been integrated into recommender systems as item-side auxiliary information, enhancing recommendation accuracy. However, constructing and integrating structural user-side knowledge remains a significant challenge due to the improper granularity and inherent scarcity of user-side features. Recent advancements in Large Language Models (LLMs) offer the potential to bridge this gap by leveraging their human behavior understanding and extensive real-world knowledge. Nevertheless, integrating LLM-generated information into recommender systems presents challenges, including the risk of noisy information and the need for additional knowledge transfer. In this paper, we propose an LLM-based user-side knowledge inference method alongside a carefully designed recommendation framework to address these challenges. Our approach employs LLMs to infer user interests based on historical behaviors, integrating this user-side information with item-side and collaborative data to construct a hybrid structure: the Collaborative Interest Knowledge Graph (CIKG). Furthermore, we propose a CIKG-based recommendation framework that includes a user interest reconstruction module and a cross-domain contrastive learning module to mitigate potential noise and facilitate knowledge transfer. We conduct extensive experiments on three real-world datasets to validate the effectiveness of our method. Our approach achieves state-of-the-art performance compared to competitive baselines, particularly for users with sparse interactions.

Code — <https://github.com/laowangzi/CIKGRec>

Introduction

In the era of information explosion, recommender systems play a crucial role in helping individuals obtain relevant information (Gao et al. 2023a; Cheng et al. 2024). However, these systems have long struggled with data sparsity and the cold start problem (Zhou et al. 2007). Recently, knowledge graphs (KGs) have been introduced as auxiliary information on the item side, enhancing recommendation algorithms’ performance and alleviating the item cold start issue (Cao et al. 2023). Technically, KGs are integrated into the user-item collaborative graph (CG), as illustrated in Fig. 1(a), to form a

hybrid structure known as the collaborative knowledge graph (CKG) (Wang et al. 2019b), depicted in Fig. 1(b). In recent years, significant progress has been made using graph neural networks (GNNs) with the CKG for recommendations (Wang et al. 2019c; Yang et al. 2022, 2023a). However, KGs provide auxiliary information solely on the item side (Wang et al. 2021), and the construction and integration of auxiliary information on the user side into recommendation algorithms remain unexplored.

Unlike item-side features, which can be naturally abstracted as structured knowledge (e.g., (movie *A*, has genre, comedy)), user-side features present significant challenges in structuring and integrating into graph-based recommendation algorithms. These challenges arise from issues related to the improper granularity and scarcity of existing user-side features. Specifically, user meta-features, such as gender, age, and nationality, are too coarse to provide precise knowledge, often leading to over-smoothing problems in GNNs. While user interests are crucial for recommendations, they are typically abstract and lack explicit user feedback (Wu et al. 2023), making them difficult to capture and utilize.

Fortunately, the rapid development of large language models (LLMs) has shown *impressive capabilities in understanding and simulating human behavior* (Tan and Jiang 2023; Hu et al. 2023b; Gao et al. 2024). Leveraging LLMs to interpret users’ historical behavior can abstract specific and meaningful user interest knowledge, thereby advancing knowledge-based recommendation algorithms. However, existing LLM-based works primarily focus on integrating recommendation tasks with dialogue systems (Geng et al. 2022; Hou et al. 2024) or using natural language features generated by LLMs as enhanced semantic information (Zhao et al. 2024; Wei et al. 2024). How to prompt LLMs to generate and construct user-side structural knowledge and combine it with knowledge-based recommendation algorithms remains an open problem.

Further incorporating LLM-generated information into recommendation algorithms introduces *two significant challenges*: (1) How to effectively utilize LLM-generated user-side knowledge while minimizing noise interference. The LLM-generated content may contain potential noise due to the hallucination problem (Zhang et al. 2023). (2) How to effectively transfer information from the auxiliary knowledge domain to the recommendation domain. Addressing these

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

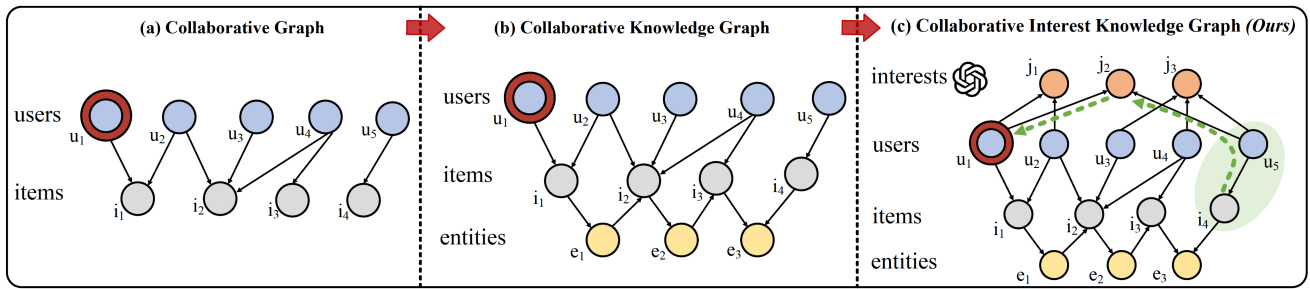


Figure 1: A toy example of the progress from traditional methods to the collaborative interest knowledge graph. u_1 is the target user to provide recommendation for. The green circle denote the important users and items discovered by user-side higher-order knowledge but is overlooked by traditional methods.

challenges is essential for developing effective downstream recommendation algorithms.

In this paper, we propose an LLM-based user-side knowledge inference method, which generates and structures user interest information into a structural knowledge format. Additionally, we introduce a recommendation framework called **CIKGRec** that integrates the generated user-side structural knowledge into a GNN-based recommendation algorithm. Specifically, we employ LLMs to infer user interest information based on users’ historical behaviors. We construct a hybrid structure, depicted in the right of Fig. 1(c), called CIKG to facilitate recommendation methods in capturing user-side information. Moreover, we propose a GNN-based recommendation algorithm to effectively utilize the CIKG information. To address the issue of additional noise introduced by LLMs, we develop an enhanced graph masked autoencoder (GMAE)-based user interest reconstruction module to improve model robustness. To address the challenge of knowledge transfer across domains, we develop a cross-domain contrastive learning module that maximizes mutual information between the auxiliary information domain and the recommendation domain. Extensive experiments on three real-world datasets demonstrate that our method outperforms baseline methods, particularly for users with sparse interactions. The main contributions of this paper are as follows:

- We propose an LLM-based user-side knowledge construction method, which enhances the existing CKG by integrating user interest knowledge, resulting in a novel hybrid structure called CIKG. This enrichment provides additional auxiliary information, which is beneficial for recommendations from a graph structure perspective.
- We further introduce a CIKG-based recommendation framework that effectively addresses the noise and information transfer challenges posed by incorporating LLM-generated user-side information into the recommendation domain.
- To verify the effectiveness of our proposed method, we conduct extensive experiments on three real-world datasets. Our method achieves optimal results overall and is particularly effective in alleviating the data sparsity problem for users.

Related Work

LLM for Recommender Systems. In recent years, the rapid development of LLMs has garnered increasing attention for their application in recommendation tasks. The role of LLMs in these tasks can be categorized into two primary functions (Fan et al. 2023; Kim et al. 2024): LLMs as recommenders and LLMs as enhancers. When used as recommenders, LLMs integrate recommendation tasks into conversational scenarios, tokenizing item IDs and user IDs. Due to their extensive real-world knowledge, LLMs as recommender exhibit advantages in few-shot scenarios (Li et al. 2024), fairness (Jiang et al. 2024) and explainability (Gao et al. 2023b; Xie et al. 2023). However, they struggle to capture the pervasive collaborative signals, resulting in limited advantages over traditional collaborative filtering algorithms (Kang et al. 2023). Some studies have explored enhancing traditional recommendation algorithms by incorporating features generated by LLMs. These models generate missing item features (Ren et al. 2024), user-item interactions (Wei et al. 2024), etc., and encode LLM-generated context from natural language into semantic vectors for enhancing traditional algorithms. Nevertheless, current methods of using LLMs for data enhancement primarily focus on the meta-features, neglecting structural knowledge from the user side. This limitation hinders their ability to facilitate traditional methods to capture higher-order collaborative information.

KG-based Recommendation Methods. Another closely related area of this work involves KG-based recommendation methods. KB4Rec (Zhao et al. 2019) links knowledge base information with items in the recommendation domain, constructing an item-side KG to facilitate the fusion of KG and the recommendation task. Existing KG-based methods can be broadly categorized into two groups: path-based methods and regularization-based methods. Path-based approaches extract higher-order information paths and input them into recommendation models (Wang et al. 2019c; Hu et al. 2018). Regularization-based techniques introduce additional loss terms to capture the KG’s structure, regulating the learning process of recommender models (Zhang et al. 2016; Wang et al. 2019b,a, 2021; Hu et al. 2023a). Recent efforts have focused on enhancing model robustness (Yang et al. 2022) and reconstructing knowledge graphs (Yang et al. 2023a; Tang et al. 2024) with information from the recommendation do-

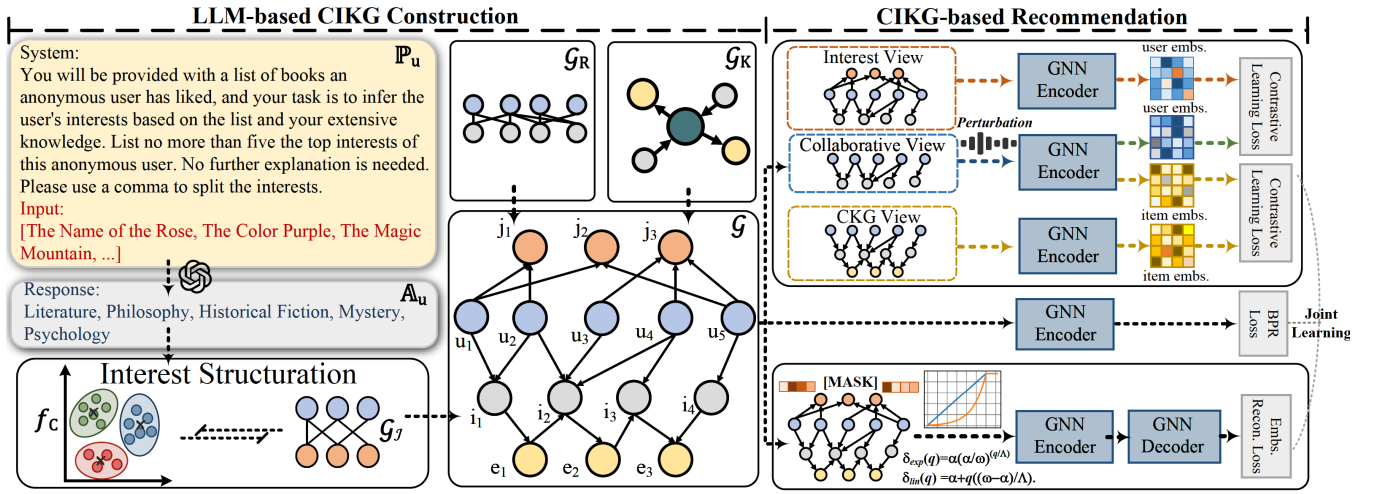


Figure 2: The overall architecture of our method. The left half illustrates the LLM-based user-side knowledge construction, and the right half shows the CIKG-based recommendation.

main. However, existing methods often overlook structured knowledge on the user side and fail to process unreliable content generated by LLMs.

Problem Definition

In this section, we first introduce the concept of the CG and the KG. Then, we formally define the top- K recommendation task.

Definition of CG. In a recommendation scenario, historical user-item interactions (e.g., watches and clicks), denoted as \mathcal{E}^+ , are represented as a user-item collaborative graph $\mathcal{G}_R = (\mathcal{V}_R, \mathcal{E}^+)$, where $\mathcal{V}_R = (\mathcal{U} \cup \mathcal{I})$, with \mathcal{U} representing the set of users and \mathcal{I} representing the set of items.

Definition of KG. In addition to interactions, item-side information is organized as a knowledge graph \mathcal{G}_K , a directed graph composed of subject-property-object triples (Shi and Weninger 2017). Formally, $\mathcal{G}_K = \{(h, r, t) | h, t \in \mathcal{V}_E, r \in \mathcal{R}\}$, where \mathcal{V}_E denotes the set of entities and \mathcal{R} denotes the set of relations. An item-entity projection set $\mathcal{P} = \{(i, e) | i \in \mathcal{I}, e \in \mathcal{V}_E\}$ links items to entities.

Top- K Recommendation. The top- K recommendation task is defined as follows: Given \mathcal{G}_R , \mathcal{G}_K and \mathcal{E}^+ , the goal is to recommend the top- K items that each user $u \in \mathcal{U}$ is most likely to interact with, among the items u has not yet interacted with.

Method

The architecture of our approach is depicted in Fig. 2. Overall, our method comprises two stages: LLM-based CIKG construction (left half of Fig. 2) and CIKG-based recommendation (right half of Fig. 2). The details of these stages are as follows.

LLM-based CIKG Construction

This section introduces the details of LLM-based user-side knowledge generation and how to convert such knowledge

from natural language into structured user-side information. Furthermore, we introduce the concept of the CIKG.

LLM-based User Interest Inference. Modeling and representing user interests is crucial for recommender systems (Gao et al. 2015; Yang et al. 2023b). Traditional recommendation algorithms represent user interests implicitly as dense vectors (Wu et al. 2023). The abstract and complex nature of user interests makes explicitly representing them challenging. This limitation hinders the recommendation model's ability to fully leverage user interest information. Recently, LLMs have shown good ability to understand user behavior (Tan and Jiang 2023; Hu et al. 2023b; Gao et al. 2024), which makes it possible to model user interests explicitly. Inspired by this, we devise prompts derived from the dataset's user behaviors to enable the LLM to infer and generate user interests not originally present in the dataset. Leveraging the LLM's extensive knowledge and reasoning ability, the model can summarize high-level, abstract user interests based on the relevance of items interacted with by users rather than relying on traditional user-side meta-features. Specific examples are shown in Fig. 2. Formally, the LLM-based user interest inference is defined as follows:

$$\mathbb{A} = \{\mathbb{A}_1, \mathbb{A}_2, \dots, \mathbb{A}_{|\mathcal{U}|}\}, \mathbb{A}_u = LLM(\mathbb{P}_u, \mathcal{E}_u^+), \quad (1)$$

where \mathbb{P}_u represents the textual prompt for user u , and \mathbb{A} represents the generated user interest textual information.

User Interest Knowledge Structuration. The purpose of structuring user interest knowledge is twofold: (1) Ensuring vector space consistency. Previous works (Ren et al. 2024; Wei et al. 2024) have encoded LLM-generated text as semantic embeddings and incorporated them into traditional recommendation methods. However, direct addition poses the problem of misalignment (Hao et al. 2019; Radford et al. 2021) between semantic signals and collaborative signals, resulting in suboptimal performance. (2) Capturing user-side multi-hop information. Explicitly modeling user interests as

structured information facilitates the model mining higher-order relationships between users with common interests. Technically, we first employ text clustering methods, which can be as simple as tf-idf or deep sentence embedding methods, to merge similar interests. In this way, we semantically avoid inconsistencies where the same interests are expressed differently due to the randomness of LLM outputs. We then connect the merged interests with users to build a user interest knowledge graph \mathcal{G}_J :

$$\begin{aligned}\mathcal{J} &= \{j_1, j_2, \dots, j_\kappa\} = f_C(\mathbb{A}, \kappa), \\ \mathcal{G}_J &= (\mathcal{V}_J, \mathcal{E}_J), \quad \mathcal{V}_J = (\mathcal{U} \cup \mathcal{J}),\end{aligned}\quad (2)$$

where $f_C(\cdot)$ denotes the textual clustering method, and κ is a hyperparameter representing the number of clusters. An edge $(u, j) \in \mathcal{E}_J$ exists if user u has the interest j .

The CIKG is defined as a unified relational graph \mathcal{G} , encoding user behaviors, user interest knowledge, and item-side knowledge. Each user behavior in \mathcal{G}_R is represented as a triplet $(u, Interact, i)$, where *Interact* denotes the relationship between user u and item i . User interest information in \mathcal{G}_J is represented as triplets $(u, Has\ Interest, j)$. By integrating the item-entity projection set \mathcal{P} , the user behaviors graph \mathcal{G}_R , and user-side information \mathcal{G}_J with the knowledge graph \mathcal{G}_K , we form the unified graph $\mathcal{G} = \{(h, r, t) | h, t \in \mathcal{V} = (\mathcal{U} \cup \mathcal{V}_E \cup \mathcal{J}), r \in \mathcal{R}'\}$, where $\mathcal{R}' = \mathcal{R} \cup \{Interact, Has\ Interest\}$, as the CIKG. Let $Z \in \mathbb{R}^{|\mathcal{V}| \times D}$ denotes the initial embedding matrix of \mathcal{G} , where D is the embedding dimension. $z_i \in \mathbb{R}^{D \times 1}$ represents the embedding of node i .

CIKG-based Recommendation

This section introduces our proposed CIKG-based recommendation method, which consists of three components: user interest reconstruction, cross-domain contrastive learning, and model training. The details are as follows.

User Interest Reconstruction. Due to the hallucination problem in LLMs (Zhang et al. 2023) and the complexity of user interests, the user interests inferred by LLMs can be inaccurate and noisy. The objective of this module is to reduce noise interference. Inspired by the robustness of graph representations achieved through GMAEs, we apply masking and reconstruction on user interest nodes, which facilitates the model’s focus on crucial information and enhances its generalization ability. Unlike previous GMAE methods that use a fixed mask rate (Hou et al. 2022) or incorporate randomness into the mask rate determination (Wang et al. 2024), we propose two mask rate acceleration scheduling strategies that dynamically control the mask rate. Drawing inspiration from curriculum learning (Wang, Chen, and Zhu 2022), these strategies progressively transform task difficulty from easy to hard. Formally, the two strategies can be expressed as:

$$p^q = \psi(\delta(q), \omega), \quad \delta(q) = \begin{cases} \delta_{lin}(q) = \alpha + q \frac{(\omega - \alpha)}{\Lambda}, \\ \delta_{exp}(q) = \alpha \left(\frac{\omega}{\alpha}\right)^{\frac{q}{\Lambda}}, \end{cases} \quad (3)$$

where $\delta(\cdot)$ denotes the scheduling function. α , ω , and Λ are hyperparameters representing the initial and maximum

mask rates, and the required epochs to convergence to the maximum mask rate. q denotes the current epoch number. p^q denotes the mask rate at epoch q . $\psi(\cdot)$ is the minimize function. It’s easy to prove that:

$$\delta_{exp}(q) < \delta_{lin}(q), 0 < q < \Lambda. \quad (4)$$

Therefore, our proposed exponential growth strategy allows the task difficulty to grow more slowly, as shown in the user interest reconstruction module of Fig. 2. After determining the mask rate p^q , we sample a subset of nodes $\tilde{\mathcal{J}} \in \mathcal{J}$ and mask their embeddings with a learnable mask token $[M]$. For each interest node j , we have:

$$z'_j = \begin{cases} z_{[M]} & \text{if } j \in \tilde{\mathcal{J}}, \\ z_j & \text{if } j \notin \tilde{\mathcal{J}}. \end{cases} \quad (5)$$

The masked embeddings Z' and \mathcal{G} are then sent to the encoder $f_G^{l-1}(\cdot)$, which is a GNN with $(l-1)$ layers. Without loss of generality, we use LightGCN (He et al. 2020) as the graph encoder for all modules. The decoder $f_G^1(\cdot)$ is then applied to obtain the reconstructed node embeddings. This process can be formulated as:

$$Z'' = f_G^{l-1}(Z', \mathcal{G}), \quad Z''' = f_G^1(\chi(Z''), \mathcal{G}), \quad (6)$$

where $\chi(\cdot)$ denotes the linear projection layer. Subsequently, we define the user interest reconstruction loss by comparing Z and Z''' . The loss function, with the scaling factor η , is described as follows:

$$\mathcal{L}_u = \frac{1}{|\tilde{\mathcal{J}}|} \sum_{j \in \tilde{\mathcal{J}}} \left(1 - \frac{z_j^T z_j'''}{\|z_j\| \cdot \|z_j'''\|}\right)^\eta, \quad \eta \geq 1, \quad (7)$$

where $\|\cdot\|$ represents the L_2 normalization function.

Cross-domain Contrastive Learning. To facilitate the transfer of user-side and item-side knowledge to the recommendation domain, we devise a cross-domain contrastive learning module. Unlike most existing contrastive learning-based models (You et al. 2020; Wu et al. 2021) that augment graphs by perturbing the graph structure and potentially losing useful information, we construct augmented views by adding auxiliary information from the user or item side to \mathcal{G}_R . Our proposed augmentation approach aligns auxiliary information with the collaborative signals of the recommendation domain, all while preserving the integrity of the graph structure. The encoded embeddings of the augmented views are as follows:

$$\begin{aligned}Z^{v1} &= f_G^l(Z, \mathcal{G}_R, \gamma), \\ Z^{v2} &= f_G^l(Z, \mathcal{G}_{v2}), \quad \mathcal{G}_{v2} = Merge(\mathcal{G}_R, \mathcal{G}_J), \\ Z^{v3} &= f_G^l(Z, \mathcal{G}_{v3}), \quad \mathcal{G}_{v3} = Merge(\mathcal{G}_R, \mathcal{M}(\mathcal{G}_K, \mathcal{P})),\end{aligned} \quad (8)$$

where γ is representation-level perturbations as depicted in SimGCL (Yu et al. 2022), which provides augmentation to \mathcal{G}_R without perturb graph structures. *Merge*(\cdot) is a function to merge the edges and nodes of graphs and eliminate co-reference nodes or edges. $\mathcal{M}(\cdot)$ is the map function to transfer entity id to item id based on \mathcal{P} .

MODEL	DBbook2014				Book-Crossing				MovieLens-1M			
	R@50	R@100	N@50	N@100	R@50	R@100	N@50	N@100	R@50	R@100	N@50	N@100
LightGCN	0.4214	0.5358	0.2230	0.2481	0.1607	<u>0.2246</u>	0.0765	0.0911	0.4086	0.5574	0.3977	0.4432
SGL	0.4258	0.5249	0.2350	0.2567	0.1613	0.2154	0.0860	0.0982	0.4094	0.5505	0.4004	0.4433
SimGCL	0.4280	0.5369	0.2376	0.2610	<u>0.1653</u>	0.2211	0.0860	0.0988	0.4131	0.5560	0.4048	0.4482
CKE	0.3419	0.4448	0.1832	0.2060	0.1284	0.1691	0.0708	0.0811	0.3948	0.5421	0.3841	0.4295
CFKG	0.3586	0.4559	0.1627	0.1816	0.1216	0.1771	0.0503	0.0617	0.3930	0.5439	0.3672	0.4142
KGAT	0.4174	0.5309	0.2110	0.2360	0.1564	0.2133	<u>0.0865</u>	<u>0.1004</u>	0.4059	0.5527	0.3928	0.4378
KGIN	0.4147	0.5337	0.2200	0.2462	0.1468	0.2075	0.0681	0.0819	0.4095	0.5581	0.4004	0.4453
KGCL	0.4308	0.5271	<u>0.2453</u>	<u>0.2666</u>	0.1562	0.2049	0.0859	0.0971	0.4083	0.5496	0.3996	0.4430
MCCLK	0.4376	0.5491	0.2054	0.2273	0.1476	0.2039	0.0656	0.0774	<u>0.4176</u>	<u>0.5655</u>	0.3876	0.4339
KGRec	<u>0.4415</u>	<u>0.5497</u>	0.2373	0.2611	0.1473	0.2075	0.0699	0.0836	0.4136	0.5633	<u>0.4072</u>	<u>0.4524</u>
CIKGRec	0.4642	0.5756	0.2494	0.2739	0.1791	0.2455	0.0889	0.1041	0.4250	0.5718	0.4162	0.4609
%Imp.	5.14%	4.71%	1.67%	2.74%	8.35%	9.31%	2.77%	3.69%	1.77%	1.11%	2.21%	1.88%

Table 1: The overall performance of CIKGRec and baseline models. The best and second-best performances are highlighted in **bold** and underline, respectively.

Dataset	DBbook2014	Book-Crossing	MovieLens-1M
#Users	5,576	6,616	6,040
#Items	2,680	8,853	3,260
#Ratings	65,961	110,662	998,539
#Density	0.44%	0.19%	5.07%
#Relations	13	4	20
#Entities	8,762	1,404	14,377
#Triplets	134,223	1,137	415,104

Table 2: Statistics of the datasets.

The positive samples of the user node in \mathcal{G}_{v_1} are the user node in \mathcal{G}_{v_2} , and the positive examples of the item node in \mathcal{G}_{v_1} are the item nodes in \mathcal{G}_{v_3} , since \mathcal{G}_{v_2} and \mathcal{G}_{v_3} add the knowledge of the user- and item-side, respectively. The contrastive objective \mathcal{L}_c , which formally maximizes agreement among positive pairs and minimizes agreement among negative pairs, is defined based on the InfoNCE loss (Chen et al. 2020) as follows:

$$\mathcal{L}_c = \sum_{u \in \mathcal{U}} -\log \frac{\exp(s(z_u^{v_1}, z_u^{v_2})/\tau)}{\sum_{w \in \mathcal{U}} \exp(s(z_u^{v_1}, z_w^{v_2})/\tau)} + \sum_{i \in \mathcal{I}} -\log \frac{\exp(s(z_i^{v_1}, z_i^{v_3})/\tau)}{\sum_{w \in \mathcal{I}} \exp(s(z_i^{v_1}, z_w^{v_3})/\tau)}, \quad (9)$$

where $s(\cdot)$ is the cosine function to estimate the similarity of positive and negative pairs. τ is the temperature coefficient to control the softness of the probability distribution.

Model Training. In addition to the two loss functions introduced above, the model training process also includes the recommendation task loss function \mathcal{L}_r , which can be formulated as:

$$\hat{Z} = f_G^l(Z, \mathcal{G}), \quad \mathcal{L}_r = \sum_{(u, i, i')} -\ln(\sigma(\hat{z}_u^T \hat{z}_i - \hat{z}_u^T \hat{z}_{i'})), \quad (10)$$

where $(u, i) \in \mathcal{E}^+$ and i' is a randomly sampled negative item of user u . σ is the sigmoid function. The main loss is the weighted sum of the three loss functions, which can be formulated as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_r + \lambda_2 \mathcal{L}_u + \lambda_3 \mathcal{L}_c, \quad (11)$$

where λ is the hyperparameter to control the measures of loss functions. In order to further enhance the multi-relational semantic representation space for entity-item dependencies, we perform the alternative translation-based training (Bordes et al. 2013), which can be formulated as:

$$\mathcal{L}_t = \sum_{(h, r, t, t')} -\ln \sigma(\|z_h + z_r - z_{t'}\| - \|z_h + z_r - z_t\|), \quad (12)$$

where $(h, r, t) \in \mathcal{G}_K$ and the negative tail t' of head h is randomly sampled from \mathcal{G}_K .

Experiments

Experimental Settings

Datasets. We conduct extensive experiments on three real-world datasets: DBbook2014 (Cao et al. 2019), Book-Crossing (Dong et al. 2017), and MovieLens-1M (Noia et al. 2016), which all contains item-side auxiliary knowledge and vary in terms of domain, size, and density. We follow the previous literature (Cao et al. 2019; Hu et al. 2024; Li et al. 2023) filtering out low-frequency users (i.e., lower than 10 in movielens and 5 in DBbook2014 and Book-Crossing). The statistics of the processed datasets are shown in Tab. 2.

Baselines. We compare the performance of CIKGRec with various baseline methods, which can be broadly categorized into three groups: (1) **GNN-based general recommenders**, including LightGCN (He et al. 2020) and two contrastive learning-based methods, SGL (Wu et al. 2021) and SimGCL (Yu et al. 2022). (2) **Embedding-based knowledge-aware recommenders**, including CKE (Zhang et al. 2016) and CFKG (Ai et al. 2018). (3) **GNN-based item-side**




Figure 3: Performance comparisons under different sparsity user groups.

knowledge-aware recommenders, including KGAT (Wang et al. 2019b), KGIN (Wang et al. 2021), and self-supervised learning-enhanced methods, KGCL (Yang et al. 2022), MC-CLK (Zou et al. 2022), KGRec (Yang et al. 2023a).

Implementation Details and Evaluation Metrics. We construct the training, validation, and test sets by randomly splitting the historical behavior of each user in a ratio of 7 : 1 : 2. For our model, we search the learning rate from 0.0001 to 0.005. We test the user interest reconstruction initial mask rate α from 0.05 to 0.15 and the maximum mask rate ω from 0.35 to 0.95. In this paper, we adopt the “gpt-3.5-turbo-0125” model as the LLM. The required convergence epoch Λ is chosen in {80, 160, 320}. For all baseline models, we utilize the source code provided by the original authors. Hyperparameters for these baselines are either adopted as recommended by the original authors or fine-tuned using the validation sets. To ensure a fair comparison, we maintain consistent common hyperparameters across all models, including an embedding dimension of 64 and an evaluation interval of 1 on all datasets. An early stopping strategy with a maximum of 2000 epochs is applied to both the baselines and our model. For all methods, we report the mean and standard deviation over 10 runs with different random seeds. To evaluate the performance of our model on the top- K recommendation task, we employ two widely used metrics: Recall@ K and NDCG@ K . These metrics respectively focus on measuring the model’s ability to retrieve relevant items and its ability to rank items. In this paper, K takes the values of 50 and 100.

Overall Performance

The overall performance of CIKGRec and baseline models is shown in Tab. 1. Based on the results, we have the following key observations: (1) Our model achieves superior performance across all metrics and datasets, demonstrating the effectiveness of our model. This also underscores the efficacy of the proposed user interest reconstruction module and cross-domain contrastive learning module. (2) The model shows notable improvements on the two data-sparse datasets (DBbook2014 and Book-Crossing), with densities of 0.44% and 0.19%, respectively. We attribute this to our proposed LLM-generated user-side structural knowledge, which effectively mitigates the data sparsity problem. (3) We observe that the knowledge-aware baselines are sensitive to the avail-

	DBbook2014		Book-Crossing		MovieLens-1M	
	R@100	N@100	R@100	N@100	R@100	N@100
Ours	0.5756	0.2739	0.2455	0.1041	0.5718	0.4609
w/o UIK	0.5673	0.2615	0.233	0.0991	0.5662	0.4527
w/o UIR	0.5701	0.2704	0.2081	0.0843	0.5713	0.4601
w/o CL	0.5732	0.2659	0.2291	0.0917	0.5583	0.4480
w/o DMR	0.5712	0.2709	0.2129	0.0853	0.5710	0.4596
w linear	0.5733	0.2716	0.2452	0.1036	0.5714	0.4603

Table 3: Ablation study of CIKGRec.

ability of item-side knowledge. Specifically, these methods perform well in datasets with abundant item-side knowledge (DBbook2014 and MovieLens-1M) but perform worse in datasets with sparse item-side knowledge (Book-Crossing). Our method performs well on both item-side knowledge-dense and knowledge-sparse datasets, highlighting the effectiveness and necessity of incorporating user-side knowledge.

User-side Data Sparsity Analysis

One motivation for exploiting user-side knowledge is to alleviate the data sparsity problem, which often limits the expressiveness of recommender systems. It is challenging to learn optimal representations for inactive users with few interactions. Following previous work (Wang et al. 2019b), we divide users into four groups based on their historical interactions in the training set. We then report the average NDCG@50 results for each user group using our model and selected baseline methods from various categories. The results are depicted in Fig. 3, where the horizontal axis represents user interaction from sparse to dense.

Fig. 3 indicates that our model demonstrates superior performance for user groups with sparse interactions across the three datasets. This highlights the advantage of our model in handling user interaction sparsity. We attribute this superiority to our proposed CIKG, which incorporates LLM-generated user-side knowledge. The introduction of user-side knowledge effectively facilitates the mining of user-side collaborative information, thereby improving performance for user groups with sparse interactions. This also demonstrates the effectiveness of our user-side knowledge construction method.

Ablation Study

To evaluate the contributions of the various modules in our model, we conducted ablation experiments while keeping the hyperparameters fixed. The results, presented in Table 3, highlight the impact of each module.

First, we observe a significant performance decrease across all three datasets in the model without user interest knowledge (UIK). This underscores the importance of the LLM-generated user-side structured knowledge. Similarly, the removal of user interest reconstruction (UIR) and contrastive learning (CL) leads to diminished performance, indicating these modules’ roles in mitigating information noise and enhancing knowledge transfer. To investigate the impact of the

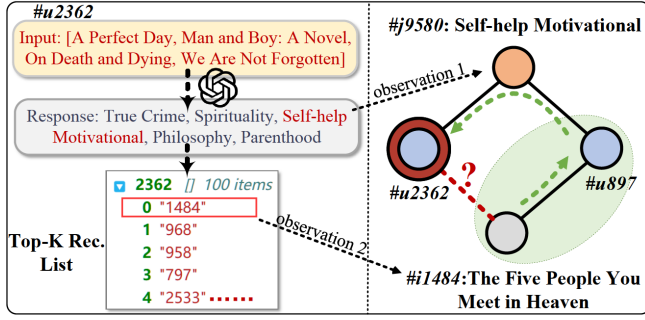


Figure 4: A real case in Book-Crossing dataset.

dynamic mask rate strategy within the UIR, we compare two variants: one without the dynamic masking rate (DMR), and one using the linear strategy. Our findings are twofold: **(1)** Both the linear and exponential growth strategies outperform the fixed masking rate, demonstrating that an easy-to-hard training approach enhances model performance. **(2)** The exponential growth strategy yields superior results compared to the linear strategy, suggesting that a slower increase in model difficulty improves performance across the datasets, validating the efficacy of our proposed exponential growth strategy.

Case Study

To intuitively assess the effectiveness of the user-side knowledge and the CIKG-based recommendation model, we select user $\#u2362$ from the Book-Crossing dataset, which exhibits the sparsest user-item interactions and item-side knowledge among the three datasets. As depicted in Fig. 4, the left panel illustrates the LLM-based interest knowledge generation process for user $\#u2362$ along with the top- K recommendation list, where higher-ranked items indicate a greater likelihood of user interaction. The right panel demonstrates the process of capturing user similarity based on the shared interest $\#j9580$ between user $\#u2362$ and user $\#u897$ using the CIKG-based recommendation model. Item $\#i1484$, a book with a positive self-improvement narrative, appears in the test set for user $\#u2362$ and in the training set for user $\#u897$.

Based on Fig. 4, we draw the following two observations: **(1)** Our LLM-based user-side knowledge method effectively summarizes the interests of user $\#u2362$, as evidenced by the alignment between the potential interact item $\#i1484$ and interest $\#j9580$ (self-help motivational). This demonstrates that our proposed LLM-based user-side knowledge inference method effectively captures user interests. **(2)** Our CIKG-based model accurately predicts user $\#u2362$'s potential interaction with item $\#i1484$, as it is prominently listed in the top- K recommendation list for this user. This result suggests that our CIKG-based recommendation model effectively leverages user-side knowledge to identify higher-order user-item similarity relationships.

Hyperparameter Sensitivity

In this section, we investigate the influence of four critical hyperparameters: the initial mask rate α , the maximum mask

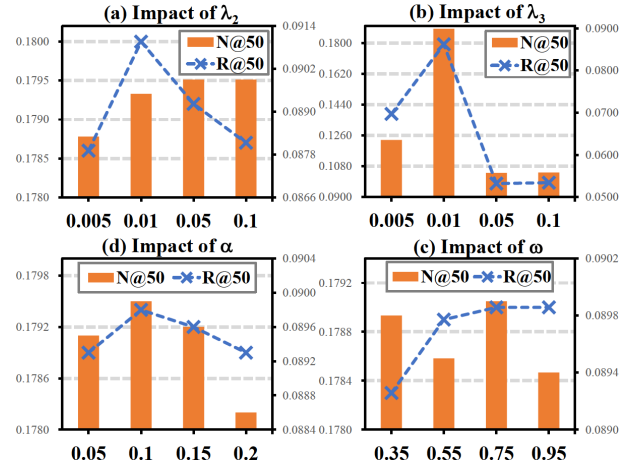


Figure 5: Hyperparameter sensitive experiments on Book-Crossing.

rate ω , the user interest reconstruction loss weight λ_2 , and the contrastive learning loss weight λ_3 . The results of the hyperparameter sensitivity experiments on the Book-Crossing dataset are illustrated in Fig.5. The primary vertical axis of the figures shows the $Recall@50$ results as a line chart, while the secondary vertical axis presents the $NDCG@50$ results as a bar chart. Based on Fig.5, we derive the following observations: **(1)** Examining λ_2 and λ_3 , which control the loss weights, reveals that model performance initially increases and then decreases as these parameters increase. Notably, a very high λ_3 leads to a significant decline in performance. This indicates that the selection of these hyperparameters should avoid extremes; they must be scaled appropriately during joint learning to achieve optimal performance. **(2)** Considering the initial mask rate α and the maximum mask rate ω , we find that a small initial mask rate, around 0.1, is beneficial for model training. The maximum mask rate ω can be quite large, even up to 0.95. We believe that these hyperparameters should be adjusted according to the specific dataset and the particular LLM employed in our model.

Conclusion

In this paper, we address the issue of missing structured knowledge on the user side, a challenge that has often been overlooked by previous knowledge-based recommender systems. We propose an LLM-based approach for generating user-side knowledge from historical user behaviors. Furthermore, we introduce a novel recommendation framework that effectively utilizes the LLM-generated user-side knowledge. This framework consists of two key modules: a user interest reconstruction module and a contrastive learning module, both of which are designed to mitigate the challenges of noise and knowledge transfer, respectively. Experimental results demonstrate that our approach effectively harnesses the knowledge embedded in LLMs and significantly enhances the performance of knowledge-aware recommendation systems.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. U24A20250 and 62176048, T2293771, 61673086, 42361144718), the Ministry of Education of Humanities and Social Science Project (No. 21JZD055), Sichuan Science and Technology Program (Grant No. 2023NSFSC1919), and the Sichuan Provincial Natural Science Foundation (Grant No. 2024NSFTD0042, 2024YFG0006, and 2024NSFSC0506). The funders had no role in the study design, data collection, analysis, decision to publish, or preparation of the manuscript.

References

- Ai, Q.; Azizi, V.; Chen, X.; and Zhang, Y. 2018. Learning Heterogeneous Knowledge Base Embeddings for Explainable Recommendation. *Algorithms*, 11(9): 137.
- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NIPS*, 2787–2795.
- Cao, Y.; Wang, X.; He, X.; Hu, Z.; and Chua, T. 2019. Unifying Knowledge Graph Learning and Recommendation: Towards a Better Understanding of User Preferences. In *WWW*, 151–161. ACM.
- Cao, Y.; Yang, L.; Wang, C.; Liu, Z.; Peng, H.; You, C.; and Yu, P. S. 2023. Multi-task Item-attribute Graph Pre-training for Strict Cold-start Item Recommendation. In *RecSys*, 322–333. ACM.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 1597–1607. PMLR.
- Cheng, Z.; Zhang, J.; Xu, X.; Trajcevski, G.; Zhong, T.; and Zhou, F. 2024. Retrieval-Augmented Hypergraph for Multimodal Social Media Popularity Prediction. In *KDD*, 445–455. ACM.
- Dong, X.; Yu, L.; Wu, Z.; Sun, Y.; Yuan, L.; and Zhang, F. 2017. A Hybrid Collaborative Filtering Model with Deep Structure for Recommender Systems. In *AAAI*, 1309–1315. AAAI Press.
- Fan, W.; Zhao, Z.; Li, J.; Liu, Y.; Mei, X.; Wang, Y.; Tang, J.; and Li, Q. 2023. Recommender Systems in the Era of Large Language Models (LLMs). *CoRR*, abs/2307.02046.
- Gao, C.; Xu, F.; Chen, X.; Wang, X.; He, X.; and Li, Y. 2024. Simulating Human Society with Large Language Model Agents: City, Social Media, and Economic System. In *WWW (Companion Volume)*, 1290–1293. ACM.
- Gao, C.; Zheng, Y.; Li, N.; Li, Y.; Qin, Y.; Piao, J.; Quan, Y.; Chang, J.; Jin, D.; He, X.; and Li, Y. 2023a. A Survey of Graph Neural Networks for Recommender Systems: Challenges, Methods, and Directions. *Trans. Recomm. Syst.*, 1(1): 1–51.
- Gao, H.; Tang, J.; Hu, X.; and Liu, H. 2015. Content-Aware Point of Interest Recommendation on Location-Based Social Networks. In *AAAI*, 1721–1727. AAAI Press.
- Gao, Y.; Sheng, T.; Xiang, Y.; Xiong, Y.; Wang, H.; and Zhang, J. 2023b. Chat-REC: Towards Interactive and Explainable LLMs-Augmented Recommender System. *CoRR*, abs/2303.14524.
- Geng, S.; Liu, S.; Fu, Z.; Ge, Y.; and Zhang, Y. 2022. Recommendation as Language Processing (RLP): A Unified Pre-train, Personalized Prompt & Predict Paradigm (P5). In *RecSys*, 299–315. ACM.
- Hao, Y.; Wang, N.; Gao, X.; Li, J.; and Wang, X. 2019. Dual-alignment Feature Embedding for Cross-modality Person Re-identification. In *ACM Multimedia*, 57–65. ACM.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*, 639–648. ACM.
- Hou, Y.; Zhang, J.; Lin, Z.; Lu, H.; Xie, R.; McAuley, J. J.; and Zhao, W. X. 2024. Large Language Models are Zero-Shot Rankers for Recommender Systems. In *ECIR (2)*, volume 14609 of *Lecture Notes in Computer Science*, 364–381. Springer.
- Hou, Z.; Liu, X.; Cen, Y.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. GraphMAE: Self-Supervised Masked Graph Autoencoders. In *KDD*, 594–604. ACM.
- Hu, B.; Shi, C.; Zhao, W. X.; and Yu, P. S. 2018. Leveraging Meta-path based Context for Top- N Recommendation with A Neural Co-Attention Model. In *KDD*, 1531–1540. ACM.
- Hu, Z.; Cai, S.; Wang, J.; and Zhou, T. 2023a. Collaborative recommendation model based on multi-modal multi-view attention network: Movie and literature cases. *Appl. Soft Comput.*, 144: 110518.
- Hu, Z.; Feng, Y.; Luu, A. T.; Hooi, B.; and Lipani, A. 2023b. Unlocking the Potential of User Feedback: Leveraging Large Language Model as User Simulators to Enhance Dialogue System. In *CIKM*, 3953–3957. ACM.
- Hu, Z.; Nakagawa, S.; Zhuang, Y.; Deng, J.; Cai, S.; Zhou, T.; and Ren, F. 2024. Hierarchical Denoising for Robust Social Recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 1–14.
- Jiang, M.; Bao, K.; Zhang, J.; Wang, W.; Yang, Z.; Feng, F.; and He, X. 2024. Item-side Fairness of Large Language Model-based Recommendation System. In *WWW*, 4717–4726. ACM.
- Kang, W.; Ni, J.; Mehta, N.; Sathiamoorthy, M.; Hong, L.; Chi, E. H.; and Cheng, D. Z. 2023. Do LLMs Understand User Preferences? Evaluating LLMs On User Rating Prediction. *CoRR*, abs/2305.06474.
- Kim, S.; Kang, H.; Choi, S.; Kim, D.; Yang, M.; and Park, C. 2024. Large Language Models meet Collaborative Filtering: An Efficient All-round LLM-based Recommender System. In *KDD*, 1395–1406. ACM.
- Li, P.; de Rijke, M.; Xue, H.; Ao, S.; Song, Y.; and Salim, F. D. 2024. Large Language Models for Next Point-of-Interest Recommendation. In *SIGIR*, 1463–1472. ACM.
- Li, Z.-T.; Nie, W.-P.; Cai, S.-M.; Zhao, Z.-D.; and Zhou, T. 2023. Exploring the topological characteristics of urban trip networks based on taxi trajectory data. *Physica A: Statistical Mechanics and its Applications*, 609: 128391.

- Noia, T. D.; Ostuni, V. C.; Tomeo, P.; and Sciascio, E. D. 2016. SPrank: Semantic Path-Based Ranking for Top- N Recommendations Using Linked Open Data. *ACM Trans. Intell. Syst. Technol.*, 8(1): 9:1–9:34.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 8748–8763. PMLR.
- Ren, Y.; Chen, Z.; Yang, X.; Li, L.; Jiang, C.; Cheng, L.; Zhang, B.; Mo, L.; and Zhou, J. 2024. Enhancing Sequential Recommenders with Augmented Knowledge from Aligned Large Language Models. In *SIGIR*, 345–354. ACM.
- Shi, B.; and Wenginger, T. 2017. ProjE: Embedding Projection for Knowledge Graph Completion. In *AAAI*, 1236–1242. AAAI Press.
- Tan, Z.; and Jiang, M. 2023. User Modeling in the Era of Large Language Models: Current Research and Future Directions. *IEEE Data Engineering Bulletin*.
- Tang, G.; Gan, X.; Wang, J.; Lu, B.; Wu, L.; Fu, L.; and Zhou, C. 2024. EditKG: Editing Knowledge Graph for Recommendation. In *SIGIR*, 112–122. ACM.
- Wang, H.; Zhao, M.; Xie, X.; Li, W.; and Guo, M. 2019a. Knowledge Graph Convolutional Networks for Recommender Systems. In *WWW*, 3307–3313. ACM.
- Wang, L.; Tao, X.; Liu, Q.; Wu, S.; and Wang, L. 2024. Rethinking Graph Masked Autoencoders through Alignment and Uniformity. In *AAAI*, 15528–15536. AAAI Press.
- Wang, X.; Chen, Y.; and Zhu, W. 2022. A Survey on Curriculum Learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(9): 4555–4576.
- Wang, X.; He, X.; Cao, Y.; Liu, M.; and Chua, T. 2019b. KGAT: Knowledge Graph Attention Network for Recommendation. In *KDD*, 950–958. ACM.
- Wang, X.; Huang, T.; Wang, D.; Yuan, Y.; Liu, Z.; He, X.; and Chua, T. 2021. Learning Intents behind Interactions with Knowledge Graph for Recommendation. In *WWW*, 878–887. ACM / IW3C2.
- Wang, X.; Wang, D.; Xu, C.; He, X.; Cao, Y.; and Chua, T. 2019c. Explainable Reasoning over Knowledge Graphs for Recommendation. In *AAAI*, 5329–5336. AAAI Press.
- Wei, W.; Ren, X.; Tang, J.; Wang, Q.; Su, L.; Cheng, S.; Wang, J.; Yin, D.; and Huang, C. 2024. LLMRec: Large Language Models with Graph Augmentation for Recommendation. In *WSDM*, 806–815. ACM.
- Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021. Self-supervised Graph Learning for Recommendation. In *SIGIR*, 726–735. ACM.
- Wu, L.; He, X.; Wang, X.; Zhang, K.; and Wang, M. 2023. A Survey on Accuracy-Oriented Neural Recommendation: From Collaborative Filtering to Information-Rich Recommendation. *IEEE Trans. Knowl. Data Eng.*, 35(5): 4425–4445.
- Xie, Z.; Singh, S.; McAuley, J.; and Majumder, B. P. 2023. Factual and informative review generation for explainable recommendation. In *AAAI*, volume 37, 13816–13824.
- Yang, Y.; Huang, C.; Xia, L.; and Huang, C. 2023a. Knowledge Graph Self-Supervised Rationalization for Recommendation. In *KDD*, 3046–3056. ACM.
- Yang, Y.; Huang, C.; Xia, L.; Huang, C.; Luo, D.; and Lin, K. 2023b. Debaised Contrastive Learning for Sequential Recommendation. In *WWW*, 1063–1073. ACM.
- Yang, Y.; Huang, C.; Xia, L.; and Li, C. 2022. Knowledge Graph Contrastive Learning for Recommendation. In *SIGIR*, 1434–1443. ACM.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In *NeurIPS*.
- Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. Are Graph Augmentations Necessary?: Simple Graph Contrastive Learning for Recommendation. In *SIGIR*, 1294–1303. ACM.
- Zhang, F.; Yuan, N. J.; Lian, D.; Xie, X.; and Ma, W. 2016. Collaborative Knowledge Base Embedding for Recommender Systems. In *KDD*, 353–362. ACM.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; Wang, L.; Luu, A. T.; Bi, W.; Shi, F.; and Shi, S. 2023. Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *CoRR*, abs/2309.01219.
- Zhao, Q.; Qian, H.; Liu, Z.; Zhang, G.; and Gu, L. 2024. Breaking the Barrier: Utilizing Large Language Models for Industrial Recommendation Systems through an Inferential Knowledge Graph. *CoRR*, abs/2402.13750.
- Zhao, W. X.; He, G.; Yang, K.; Dou, H.; Huang, J.; Ouyang, S.; and Wen, J. 2019. KB4Rec: A Data Set for Linking Knowledge Bases with Recommender Systems. *Data Intell.*, 1(2): 121–136.
- Zhou, T.; Ren, J.; Medo, M.; and Zhang, Y.-C. 2007. Bipartite network projection and personal recommendation. *Physical review E*, 76(4): 046115.
- Zou, D.; Wei, W.; Mao, X.; Wang, Z.; Qiu, M.; Zhu, F.; and Cao, X. 2022. Multi-level Cross-view Contrastive Learning for Knowledge-aware Recommender System. In *SIGIR*, 1358–1368. ACM.