DisCo: Graph-Based Disentangled Contrastive Learning for Cold-Start Cross-Domain Recommendation

Hourun Li[*1,2], **Yifan Wang**[*3], **Zhiping Xiao**[4†], **Jia Yang**[2], **Changling Zhou**[2], **Ming Zhang**[1†], **Wei Ju**[5†]

[1] State Key Laboratory for Multimedia Information Processing, School of Computer Science, PKU-Anker LLM Lab, Peking University, Beijing, China
[2] Computer Center, Peking University, Beijing, China
[3] School of Information Technology & Management, University of International Business and Economics, Beijing, China
[4] Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA, USA
[5] College of Computer Science, Sichuan University, Chengdu, China
{lihourun, yangj, zclfly, mzhang_cs}@pku.edu.cn,
yifanwang@uibe.edu.cn, patxiao@uw.edu, juwei@scu.edu.cn

Abstract

Recommender systems are widely used in various real-world applications, but they often encounter the persistent challenge of the user cold-start problem. Cross-domain recommendation (CDR), which leverages user interactions from one domain to improve prediction performance in another, has emerged as a promising solution. However, users with similar preferences in the source domain may exhibit different interests in the target domain. Therefore, directly transferring embeddings may introduce irrelevant source-domain collaborative information. In this paper, we propose a novel graph-based disentangled contrastive learning framework to capture fine-grained user intent and filter out irrelevant collaborative information, thereby avoiding negative transfer. Specifically, for each domain, we use a multi-channel graph encoder to capture diverse user intents. We then construct the affinity graph in the embedding space and perform multi-step random walks to capture high-order user similarity relationships. Treating one domain as the target, we propose a disentangled intent-wise contrastive learning approach, guided by user similarity, to refine the bridging of user intents across domains. Extensive experiments on four benchmark CDR datasets demonstrate that DisCo consistently outperforms existing state-of-the-art baselines, thereby validating the effectiveness of both DisCo and its components.

Introduction

The rapid expansion of Internet services has deeply integrated online platforms into our daily lives, resulting in an enormous increase in digital information (Zhu et al. 2016; Zhang, Liu, and Wang 2020; Wang et al. 2022d,b; Ju et al. 2022; Qu et al. 2024; Ju et al. 2024a). Consequently, personalized recommender systems are crucial for guiding users through extensive options to identify items that align with their preferences. Collaborative Filtering (CF), which models relationships between users and dependencies among items, has seen significant success in recommendation systems. However, these CF-based methods persistently suffer from the cold-start problem, where new users lack sufficient observed interactions, making it challenging to learn effective representations for recommendations.

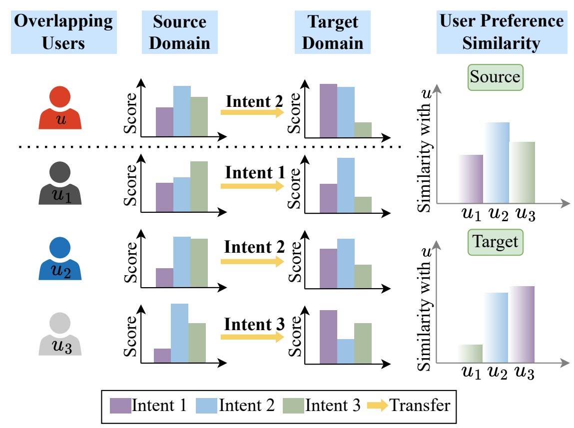

Figure 1: An illustration of user similarity distortion in CDR. Users may share diverse intents across domains. As a result, for the target user u, some users (i.e., u_1, u_3) may have similar preferences in the source domain but differ significantly in the target domain, leading to negative transfer.

To alleviate the above issue, cross-domain recommendation (CDR) has garnered considerable attention in recent years. The main idea of CDR methods is to leverage user interaction data from related domains to improve prediction accuracy in the target domain. A particularly challenging aspect is cold-start CDR, where users have interactions with items in one domain but not in the other. Traditional Embedding and Mapping (EMCDR) paradigms (Man et al. 2017; Kang et al. 2019; Salah, Tran, and Lauw 2021) encode user preferences in two domains separately and learn a mapping function from the source to the target domain, which often overlooks the diverse user-specific preferences (Wang et al.

[*]These authors contributed equally.
[†]Corresponding authors.

2022a). Additionally, some meta-learning approaches (Zhu et al. 2021, 2022; Guan et al. 2022; Li et al. 2024b) treat different user CDR as individual tasks to achieve user-specific preference transfer, where the transformed preference can be utilized as the initial embedding for the cold-start users.

Nevertheless, it still remains challenging to address the negative transfer issue for CDR (Li et al. 2024a). Since information from the source domain is not always relevant to the target domain, indiscriminately incorporating source domain data during training may lead to negative effects. Recent efforts attempt to address this problem by designing disentangled user representations that transfer only the relevant information. For example, CDRIB (Cao et al. 2022b) introduces two information bottleneck regularizers to encourage the learned representation to encode domain-shared information while limiting domain-specific information. DisenCDR (Cao et al. 2022a) learns disentangled representations to separate domain-shared and domain-specific user preferences. UniCDR (Cao et al. 2023) proposes a unified framework to capture domain-shared and domain-specific user preferences for different CDR scenarios.

Despite their promising performance, we find that most of the cold-start CDR methods focus on the user-embedding transfer, neglecting the more detailed underlying intent and preference correlations among users. In fact, the complex formation of user preferences requires to infer the latent intent under implicit interaction. Meanwhile, users who exhibit similar preferences in the source domain may have different interests in the target domain. Therefore, collaborative knowledge from the source domain can become irrelevant or even noisy in the target domain, ending up counter-effective. We demonstrate such examples in Figure 1, where user preferences are fine-grained with individual intents. Suppose we aim to learn the user preference of u in the target domain. Given the preferences of u_1, the similarity in the source domain may not hold in the target domain. As a result, learning from the irrelevant preferences of these known users can introduce bias and induce sub-optimal performance.

Motivated by the above observation, we propose **DisCo**, a novel graph-based Disentangled Contrastive learning framework for cold-start CDR, which is capable of capturing multiple intents of users and filtering our irrelevant source domain collaborative information. Specifically, we first utilize a multi-channel graph encoder to discern the underlying intent of users. Then, given the affinity graph calculated in the embedding space of two domains, we perform the multi-step random walks for each anchor user to obtain the high-order user similarities of each domain. Besides, we treat one of the domains as the target domain and propose an intent-wise contrastive architecture that performs both intra-domain and inter-domain contrastive learning to retain user similarity information for the target domain with the help of a cross-domain decoder. In this way, when introducing the source domain information, target domain-specific user preferences could be preserved while irrelevant collaborative information could be effectively filtered out by regularizing the rationale that bridging two domains.

In this paper, we make the following contributions:

- *Conceptual:* We highlight the negative transfer problem in cold-start CDR and propose to learn the more detailed disentangled user intent representation enhanced by the interaction graph to filter out irrelevant information.
- *Methodological:* We propose a novel intent-wise contrastive learning framework to retrain the user similarity information of the target domain, which could preserve the target domain-specific user preference and explicitly exploit the rationale that bridging two domains.
- *Experimental:* We conduct extensive experiments on various public datasets to evaluate DisCo. Experimental results demonstrate the superiority of our proposed framework for the cold-start CDR task. The code is released on https://github.com/HourunLi/2025-AAAI-DisCo

Related Work

Cross-Domain Recommendation

Recent CDR efforts can be categorized into two types. Intra-domain CDRs (Singh and Gordon 2008; Li, Yang, and Xue 2009; Hu, Zhang, and Yang 2018; Li and Tuzhilin 2020; Liu et al. 2020) address data sparsity by transferring abundant information from other domains for domains with limited user interactions. CBT (Li, Yang, and Xue 2009) introduces a cluster-level pattern matrix to transfer the rating patterns. CMF (Singh and Gordon 2008) adapts the matrix decomposition to jointly factorize rating matrices across domains and shares the user latent factors. CoNet (Hu, Zhang, and Yang 2018), DDTCDR (Li and Tuzhilin 2020), and Bi-TGCF (Liu et al. 2020) leverage deep models with information transfer modules. In contrast, inter-domain CDRs, or cold-start CDRs, tackle the more challenging task of recommending items to cold-start (non-overlapping) users with no prior interactions in a target domain. EMCDR (Man et al. 2017) pre-trains user embeddings of each domain and maps them by overlapping users. SSCDR (Kang et al. 2019) and SA-VAE (Salah, Tran, and Lauw 2021) extend mapping with semi-supervised learning and variational autoencoder, respectively. PTUPCDR (Zhu et al. 2022) and TMCDR (Zhu et al. 2021) treat different user CDRs as individual tasks and use personalized meta-networks for user-specific preference transfer. UniCDR (Cao et al. 2023) unifies intra- and inter-CDR in a single framework. However, most methods focus on user embedding transfer, neglecting underlying intent and the domains-specific user preferences.

Disentangled Representation Learning

Disentangled representation learning seeks to develop factorized representations that can effectively distinguish and separate the underlying explanatory factors within observed data (Bengio, Courville, and Vincent 2013). Existing researches have primarily focused on computer vision (Higgins et al. 2016; Chen et al. 2016), natural language processing (Cheng et al. 2020; Wang et al. 2022e), and graph learning (Ma et al. 2019a; Li et al. 2021; Wang et al. 2020b, 2024; Ju et al. 2024d). For example, InfoGAN (Chen et al. 2016) separates representation into noise and an additional class code, estimating mutual information (MI) between the class code and corresponding data for controllable image generation. DGCL (Li et al. 2021) introduces a self-supervised factor-wise contrastive learning framework to learn disentangled graph representation. Recently, there has been a notable surge of interest in applying disentangled representation learning techniques to recommendation systems (Ma et al. 2019b; Wang et al. 2020a, 2022c; Qin et al. 2023). MacridVAE (Ma et al. 2019b) analyzes user behavior data via disentangling representations into macro and micro levels to investigate hierarchical user intentions. DisenHAN (Wang et al. 2020a) iteratively identifies the dominant aspects of various relations within a Heterogeneous Information Network (HIN) for the recommendation. For CDR task, DR-MTCDR (Guo et al. 2023), CDRIB (Cao et al. 2022b), DisenCDR (Cao et al. 2022a) and GDCCDR (Liu et al. 2024) separate user preferences into domain-shared and domain-specific parts, focusing on identifying and transferring the shared aspects to improve performance. In this paper, we propose a disentangled learning framework to capture more detailed user intent and filter out irrelevant collaborative information to avoid negative transfer.

Notations and Problem Definition

Notations. Given the two domains $\mathcal{S}$ and $\mathcal{T}$, let $\mathcal{D}_\mathcal{S} = \{\mathcal{U}_\mathcal{S}, \mathcal{V}_\mathcal{S}, \mathcal{E}_\mathcal{S}\}$ and $\mathcal{D}_\mathcal{T} = \{\mathcal{U}_\mathcal{T}, \mathcal{V}_\mathcal{T}, \mathcal{E}_\mathcal{T}\}$ denote the corresponding data of source and target domains, where $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{E}$ denote the user, item and interaction set in the domain. The binary interaction matrix can be represented as $Y = \{y_{uv}\} \in \mathbb{R}^{|\mathcal{U}|\times|\mathcal{V}|}$, where $y_{uv} = 1$ if $(u, v) \in \mathcal{E}$, otherwise, $y_{uv} = 0$. Given the interaction matrix of source domain $Y_\mathcal{S}$ and target domain $Y_\mathcal{T}$, we construct the bipartite graphs $\mathcal{G}_\mathcal{S}$ and $\mathcal{G}_\mathcal{T}$ to depict the user-item relations of two domains. In particular, the user sets $\mathcal{U}_\mathcal{S}$ and $\mathcal{U}_\mathcal{T}$ contain an overlapping user subset $\mathcal{U}_o$. Then, the user set can be formulated as $\mathcal{U}_\mathcal{S} = \{\mathcal{U}_s, \mathcal{U}_o\}$, $\mathcal{U}_\mathcal{T} = \{\mathcal{U}_t, \mathcal{U}_o\}$, where $\mathcal{U}_s$ and $\mathcal{U}_t$ are the non-overlapping user set in each domain.

Problem Definition. Given the observed data $\mathcal{D}_\mathcal{S}$ and $\mathcal{D}_\mathcal{T}$ from the source and target domains, cold-start CDR aims to make recommendations in the target domain for non-overlapping users who are only observed in the source domain. Formally, given a user $u \in \mathcal{U}_s$ from the source domain, we seek to learn disentangled representations of user u that capture diverse intents, in order to learn the prediction function $\hat{y}_{uv} = \mathcal{F}(u, v|\Theta, \mathcal{G}_\mathcal{S}, \mathcal{G}_\mathcal{T})$, where item $v \in \mathcal{V}_\mathcal{T}$ is from the target domain. We use $\hat{y}_{uv}$ to denote the probability of user u engaging with item v and use Θ to denote the set of model parameters for the prediction function $\mathcal{F}$.

The Proposed Framework

Overview

The fundamental concept of our framework is to alleviate the negative transfer problem when introducing the source domain information for cold-start CDR. As shown in Figure 2, there are three components in our framework. Given the interaction data $\mathcal{D}_\mathcal{S}$ and $\mathcal{D}_\mathcal{T}$ from the source and target domains, we construct bipartite graph $\mathcal{G}_\mathcal{S}$ and $\mathcal{G}_\mathcal{T}$ and extract fine-grained user intents for each domain via the proposed graph encoder. Next, affinity graphs are constructed within the corresponding embedding spaces of both domains. By applying multi-step random walks based on affinity graphs, we progressively obtain user similarities as the collaborative information in both the source and target domain. And a disentangled intent-wise contrastive learning framework is proposed to retain the user similarity relationships while explicitly identifying the rationale between two domains to avoid the negative transfer. Finally, the disentangled user intents of the cold-start user are used to predict matching scores for cross-domain recommendations.

Disentangled User Intent Graph Encoder

As graph neural networks (GNNs) show a strong ability to model user interactions for the recommender (Ju et al. 2024b,c), we leverage GNNs as the backbone to capture the fine-grained user intents in each domain. For the constructed graph from the source and target domain, denoted as $\mathcal{G}_\mathcal{S}$ and $\mathcal{G}_\mathcal{T}$, we iteratively update the representations of user and item nodes through a message-passing mechanism. Taking user u as an example, the $\text{GNN}^l(\cdot)$ can be:

$$\begin{aligned} z_u^{l+1} &= \sigma\Big(W_1^l z_u^l + \sum_{v\in\mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u||\mathcal{N}_v|}} m_{uv}\Big), \\ m_{uv} &= W_1^l z_v^l + W_2^l(z_v^l \odot z_u^l), \end{aligned} \tag{1}$$

where z_u^l denotes the updated user embedding at l-th layer, m_{uv} represents the neighbor messages from the interaction item set $\mathcal{N}_u$, $\sigma(\cdot)$ is the activate function (i.e., LeakyReLU) and $W_{1/2}$ are trainable parameters. After applying L traditional message-passing layers, we obtain the extracted embedding matrix Z^L. We further utilize a graph-disentangled layer to extract the more detailed user intent with a separate channel. For each channel, we adopt GNN_k^{L+1} and the disentangled user/item representation can be obtained as:

$$Z_k = \text{GNN}_k^{L+1}(Z^L, Y), \tag{2}$$

where $k \in [1, K]$ and K represents the number of intents.

Disentangled Intent-wise Contrastive Learning

To tackle the differences in collaborative information between domains, we propose a disentangled intent-wise contrastive learning framework that uses user similarity to capture fine-grained user intents and transfer relevant information from the source domain to the target domain.

Intra-Domain Contrast. For each domain, we introduce two Siamese encoders, an online and a momentum-based target encoder, to generate effective self-supervised signals for fine-grained user intent learning. Specifically, both the online and target encoders share the same architecture and process the interaction data of each domain to obtain the user intents, denoted as Z_k and $\hat{Z}_k$. The target encoder is updated using the Exponential Moving Average (EMA) of the online encoder. Then, the user similarity in a domain for the batch instances under intent k can be defined as:

$$R_{ij} = \exp(-\|\hat{z}_{i,k} - \hat{z}_{j,k}\|/\tau), \tag{3}$$

where $\hat{z}_{i,k}$ represents the k-th channel of the output target embedding for user $u_i \in \mathcal{U}$. By treating user similarity as

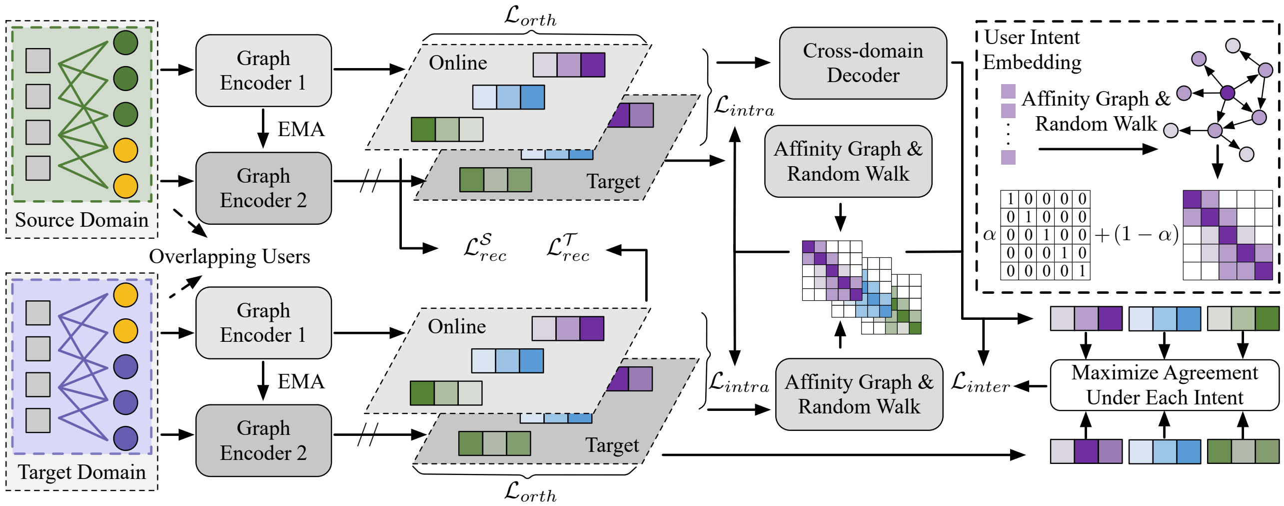

Figure 2: Illustration of the proposed framework DisCo.

the edge weight between users, a fully connected user affinity graph R could be constructed. And we normalize it in a row-wise manner to obtain the random walk transition matrix $\tilde{R}$. Considering a d step random walks (Lu et al. 2024), the high-order user similarity under intent k could be:

$$T_k = \alpha I + (1-\alpha)\tilde{R}^d, \tag{4}$$

where α is a trade-off parameter between self and user pair similarities. We take the similarity as the pseudo target and formalize the intra-domain contrastive loss as:

$$\mathcal{L}_{intra} = \sum_{u=1}^{|\mathcal{U}|}\sum_{k=1}^{K} H(T_k, \rho(Z_k, \hat{Z}_k)), \tag{5}$$

where $H(\cdot,\cdot)$ denotes the cross entropy, $\rho(\cdot,\cdot)$ is the pairwise similarity with the row-wise normalization manner:

$$[\rho(Z_k, \hat{Z}_k)]_{ij} = \frac{\exp(\phi(z_{i,k}, \hat{z}_{j,k})/\tau)}{\sum_{j'=1}^{|\mathcal{U}|}\exp(\phi(z_{i,k}, \hat{z}_{j',k})/\tau)}, \tag{6}$$

where τ is the temperature parameter, and $\phi(\cdot,\cdot)$ denotes the similarity function. Moreover, we aim to learn disentangled user representation for diverse user intents, which means $z_{i,k}$ and $z_{i,k'}$ under channel k and k' are independent of each other. Since orthogonality is a specific instance of linear independence, we loose the representation constraint to orthogonality (as Equation 7), a method has been demonstrated to be effective in several previous studies (Liang, Li, and Madden 2020; Wang et al. 2024).

$$\mathcal{L}_{orth} = |Z_k^{\mathrm{T}} Z_k - I| + |\hat{Z}_k^{\mathrm{T}}\hat{Z}_k - I|. \tag{7}$$

Inter-Domain Contrast. Given the extracted fine-grained user intents, we aim to transfer cross-domain information while maintaining distinct user collaborative information for each domain to highlight domain differences. Thus, simply maximizing the consistency between domains may lead to a loss of domain-specific information. To solve this issue, we introduce a user representation decoder that could preserve the domain-specific information while capturing the domain-shared information between the source and target domains. To be specific, for an overlapping user $u_i \in \mathcal{U}_o$, we project the source domain embedding $z_{i,k}^s$ to learn its corresponding target domain embedding.

$$e_{i,k}^{(s\to t)} = g^{(s\to t)}(z_{i,k}^s), \tag{8}$$

where $g(\cdot)$ denotes the decoder, and we simply implemented it with an MLP. To maintain user collaborative information in the target domain and prevent it from being disrupted by the incoming source domain information, we compute the similarity among overlapping users in the target domain by integrating all intents, namely, $T^t = 1/K\sum_{k=1}^{K} T_k^t$. In this way, we use it as a supervision signal and the inter-domain contrastive loss can be defined as:

$$\mathcal{L}_{inter} = \sum_{u_i=1}^{|\mathcal{U}_o|} H(T^t, p(u_j|u_i)) = -\sum_{u_i=1}^{|\mathcal{U}_o|} T_{ij}^t \log p(u_j|u_i) \tag{9}$$

where $p(u_j|u_i)$ denotes the user similarity identification task. Considering the latent fine-grained user intents, the cross entropy objective can be defined as:

$$H(T^t, p(u_j|u_i)) = -T_{ij}^t \log \mathbb{E}_{p(k|u_i)} p(u_j|u_i, k), \tag{10}$$

where $p(k|u_i)$ is the prior distribution over latent intents for user u_i. We introduce K intent prototypes $\{c_k\}_{k=1}^{k}$, and the prior distribution can be obtained as:

$$p(k|u_i) = \frac{\exp(\phi(e_{i,k}^{s\to t}, c_k))}{\sum_{k'=1}^{K}\exp(\phi(e_{i,k'}^{s\to t}, c_{k'}))}. \tag{11}$$

$p(u_j|u_i, k)$ denotes the user similarity under intent k. However, directly minimizing the objective is difficult because of latent user intents. Instead, we adopt an EM algorithm to solve the problem, where the posterior distribution $p(k|u_i, u_j)$ is defined with Bayes' theorem:

$$p(k|u_i, u_j) = \frac{p(k|u_i)p(u_j|u_i, k)}{\sum_{k=1}^{K} p(k|u_i)p(u_j|u_i, k)}. \tag{12}$$

$p(k|u_i, u_j)$ reflects how well the k-th intent aligns u_i and u_j in source and target domains. However, we cannot compute the posterior distribution because calculating $p(u_j|u_i, k)$ would require considering all users in the dataset. Therefore, we maximize the evidence lower bound (ELBO) of the log-likelihood, which can be defined as:

$$\begin{aligned}\log p(u_j|u_i) \geq & \mathbb{E}_{q(k|u_j,u_i)}[\log p(u_j|u_i, k)] \\ & - D_{KL}(q(k|u_j, u_i)\|p(k|u_i)),\end{aligned} \tag{13}$$

where $q(k|u_j, u_i)$ is a variational distribution used to approximate the posterior distribution, calculated as follows:

$$q(k|u_j, u_i) = \frac{p(k|u_i)\hat{p}(u_j|u_i, k)}{\sum_{k=1}^{K} p(k|u_i)\hat{p}(u_j|u_i, k)}, \tag{14}$$

where $\hat{p}(u_j|u_i, k)$ is the user similarity identification under intent k, calculated within a mini-batch $\mathcal{B} \subset \mathcal{U}_o$:

$$\hat{p}(u_j|u_i, k) = \frac{\exp(\phi(e_{i,k}^{s\to t}, \hat{z}_{j,k}^t))}{\sum_{j'\in\mathcal{B}}^{|\mathcal{B}|} \exp(\phi(e_{i,k}^{s\to t}, \hat{z}_{j',k}^t))}. \tag{15}$$

Notice the optimizing process is a variation of Variational EM algorithm, where we infer $q(k|u_j, u_i)$ at the E-step and optimize the ELBO at the M-step.

User Intent Adaptation and Prediction

For any single domain, we obtain the disentangled item representation $z_v = \{z_{v,1}, \ldots, z_{v,K}\}$ and the disentangled user representation $z_u = \{z_{u,1}, \ldots, z_{u,K}\}$, which represent the diverse K user intents. We use the inner product of user and item representations for each intent as the predictive function to estimate the likelihood of their interaction under that specific intent. Thereafter, we weight the sum of all intents to obtain the overall user preference,

$$r_{uv} = \sum_{k=1}^{K} p(k|u) \cdot z_{u,k}^{\mathrm{T}} z_{v,k}. \tag{16}$$

We constrain the user preference score in the range of $[0, 1]$ as the final matching score, which can be computed using a probabilistic function like the logistic function:

$$\hat{y}_{uv} = \mathrm{sigmoid}(r_{uv}) = \frac{1}{1 + \exp(r_{uv})}. \tag{17}$$

Following the previous work (Cao et al. 2023), we use the binary cross entropy (BCE) loss and train the model using a negative sampling strategy. Taking the target domain $\mathcal{T}$ as an example, the recommendation loss is defined as follows:

$$\mathcal{L}_{rec}^{\mathcal{T}} = -\sum_{(u,v)\in\mathcal{D}_{\mathcal{T}}} \log \hat{y}_{uv} - \sum_{(u,v^-)\in\mathcal{D}_{\mathcal{T}}^-} \log(1 - \hat{y}_{uv^-}), \tag{18}$$

where $\mathcal{D}_{\mathcal{T}}^-$ denotes the negative pairs uniformly sampled from unobserved interactions in the target domain.

Objective Function

The final loss term consists of two parts: 1) the contrastive loss, including the intra-domain contrastive loss $\mathcal{L}_{intra}$, the orthogonality loss $\mathcal{L}_{orth}$ and the inter-domain contrastive loss $\mathcal{L}_{inter}$. 2) the source and target domain prediction loss $\mathcal{L}_{rec}^{\mathcal{S}}$ and $\mathcal{L}_{rec}^{\mathcal{T}}$. And the overall objective is defined as:

$$\begin{aligned}&\mathcal{L}_{rec} = \mathcal{L}_{rec}^{\mathcal{S}} + \mathcal{L}_{rec}^{\mathcal{T}}, \\ &\mathcal{L}_{contra} = \beta\mathcal{L}_{inter} + (1-\beta)(\mathcal{L}_{intra} + \gamma\mathcal{L}_{orth}), \\ &\mathcal{L}_{total} = \lambda\mathcal{L}_{contra} + (1-\lambda)\mathcal{L}_{rec}.\end{aligned} \tag{19}$$

where β, γ and λ are hyper-parameters that control the weights of each part of losses. We optimize the overall objective in a mini-batches manner.

Experiment

Experimental Setup

Datasets. We experiment on four domain pairs from the public Amazon dataset [1], namely music-movie, phone-electronic, cloth-sport, and game-video. We filter out the users and items with fewer than 5 and 10 interactions in their respective domains to improve data quality. Following prior works (Cao et al. 2022b, 2023), we randomly select 20% of overlapping users (i.e., those observed in both source and target domains) and treat them as cold-start users by removing their target domain interactions during testing and validation, using the remaining users for training.

Baselines. We compare DisCo with two baseline groups: (A) Single-domain recommendation, where interactions from both domains are merged into one and standard CF-based methods are applied, including CML (Hsieh et al. 2017), BPR-MF (Rendle et al. 2012), and NGCF (Wang et al. 2019). (B) Cross-domain recommendation (CDR) models follow two paradigms. The first, Embedding and Mapping approach for CDR (EMCDR) (Man et al. 2017) pre-trains single-domain models (e.g., CML, BPR-MF, NGCF) to initialize user/item representations and includes extensions like SSCDR (Kang et al. 2019), TMCDR (Zhu et al. 2021), and SA-VAE (Salah, Tran, and Lauw 2021). The second, disentangled representation learning encodes domain-shared and domain-specific knowledge, as seen in CDRIB (Cao et al. 2022b) and UniCDR (Cao et al. 2023).

Evaluation. We employ the leave-one-out approach to evaluate the performance of all methods. Given a ground truth interaction (u, v) in target domain $\mathcal{T}$, we randomly select 999 items from the item set $\mathcal{V}^{\mathcal{T}}$ as negative samples. We then generate 1,000 records (1 positive and 999 negative samples) using the learned representation z_u^s from source domain $\mathcal{S}$ and randomly-selected positive sample z_v^t or negative sample $z_{v'}^t$ from the target domain $\mathcal{T}$. We rank the record list and use two widely-used recommendation evaluation metrics, namely NDCG@10 and HR@10, to evaluate the top-10 recommendation performance.

Implement Details. In our experiments, we set the embedding dimension to 128, the batch size to 1,024, and the slope of LeakyReLU to 0.05. In specific, we tune the number of the graph encoder layer L in the range of $[1, 6]$, the number of latent factors K in the range of $[1, 6]$, and the values of β and λ of the objective function in the range of $[0, 0.5]$. Additionally, we turn the dropout rate in the range $[0, 0.5]$, and

[1] http://jmcauley.ucsd.edu/data/amazon/index_2014.html

Dataset	Metric@10	Single-Domain Methods			Cross-domain Methods						Ours
		CML	BPRMF	NGCF	EMCDR	SSCDR	TMCDR	SA-VAE	CDRIB	UniCDR	DisCo
Sport	HR	5.82±0.20	5.75±0.26	7.22±0.11	7.41±0.16	7.27±0.02	7.18±0.07	7.51±0.02	**11.10±0.29**	10.67±0.19	10.72±0.32
	NDCG	3.29±0.16	3.16±0.15	3.63±0.07	4.03±0.12	3.75±0.02	3.84±0.04	3.72±0.02	5.78±0.15	**6.19±0.17**	5.81±0.26
Cloth	HR	6.97±0.11	6.75±0.13	7.07±0.30	7.91±0.15	6.12±0.05	8.11±0.16	7.21±0.05	12.23±0.24	12.28±0.24	**12.85±0.42**
	NDCG	3.92±0.14	3.26±0.15	3.48±0.13	5.17±0.08	3.06±0.04	5.05±0.12	4.59±0.08	6.79±0.22	**7.31±0.26**	6.92±0.32
Game	HR	2.82±0.18	3.77±0.40	5.14±0.22	5.07±0.17	3.48±0.06	5.36±0.09	5.84±0.13	8.72±0.41	7.68±0.32	**9.54±0.28**
	NDCG	1.44±0.09	1.89±0.19	2.73±0.09	2.44±0.08	1.59±0.03	2.58±0.07	2.78±0.06	4.58±0.20	4.63±0.17	**4.87±0.27**
Video	HR	3.07±0.10	4.46±0.56	7.41±0.18	8.43±0.04	5.51±0.08	8.85±0.11	7.46±0.13	12.66±0.39	10.32±0.36	**13.38±0.31**
	NDCG	1.30±0.08	2.36±0.34	3.87±0.10	4.29±0.02	2.61±0.02	4.41±0.08	3.71±0.06	6.66±0.20	5.43±0.16	**6.86±0.14**
Music	HR	8.70±0.27	8.74±0.13	8.86±0.10	8.95±0.12	3.59±0.04	9.48±0.08	8.57±0.15	14.72±0.23	11.59±0.35	**15.92±0.21**
	NDCG	4.53±0.14	4.72±0.04	4.15±0.09	4.70±0.08	1.82±0.01	5.15±0.05	4.48±0.06	7.98±0.10	6.08±0.22	**8.56±0.13**
Movie	HR	7.87±0.11	9.12±0.15	9.92±0.16	11.89±0.10	5.47±0.03	10.22±0.15	11.51±0.18	14.86±0.44	12.40±0.45	**16.33±0.28**
	NDCG	3.95±0.03	4.72±0.08	4.53±0.10	6.06±0.14	2.72±0.01	6.00±0.04	5.66±0.11	7.63±0.23	6.41±0.22	**8.39±0.17**
Phone	HR	11.56±0.23	11.98±0.22	13.62±0.20	13.84±0.07	6.14±0.03	13.95±0.22	14.35±0.29	18.37±0.32	14.27±0.37	**18.74±0.31**
	NDCG	6.41±0.13	7.65±0.07	7.47±0.12	7.73±0.03	3.33±0.01	7.56±0.09	7.66±0.18	10.16±0.25	8.87±0.25	**10.19±0.22**
Elec	HR	12.44±0.10	12.78±0.33	16.24±0.21	17.01±0.03	10.34±0.03	16.38±0.08	17.21±0.19	20.96±0.31	15.67±0.47	**21.03±0.29**
	NDCG	6.85±0.12	7.46±0.23	9.09±0.11	9.78±0.02	5.50±0.01	9.24±0.05	9.49±0.13	**12.01±0.22**	9.25±0.30	11.54±0.23

Table 1: Performance comparison (expressed in %) of CDR on cross domain recommendations. The best performance is **bold-faced** and the runner-up is underlined in terms of the corresponding metric.

the learning rate in the range of $[0, 0.005]$. For all methods, we tune the hyper-parameters by grid search, run each experiment five times with different random seeds and record the best result of each time according to the HR@10 performance on the validation set.

Performance Analysis

Table 1 presents the mean best recommendation performance for all methods evaluated across four CDR scenarios, and we have several key observations: (1) Graph-based approaches like NGCF consistently outperform methods CML and BPR-MF, highlighting the effectiveness of graph convolution in capturing multi-hop neighborhood and high-order interaction information. (2) Compared to the single-domain approaches, cross-domain methods show a superior performance, which demonstrates the ability of the embedding and mapping paradigm to capture domain-specific knowledge and nuances between source and target domains. (3) Compared to cross-domain approaches, DisCo excels on the game-video and music-movie datasets and remains competitive on the sport-cloth and phone-electronic datasets when compared to state-of-the-art graph-encoder models such as CDRIB and UniCDR. Additionally, DisCo outperforms all other EMCDR-based methods, indicating the effectiveness of intra-domain and inter-domain contrast mechanisms to filter out irrelevant collaborative information.

Ablation Study

To evaluate the effectiveness of the various modules in DisCo and understand their functions, we compare DisCo with four variants: (1) Variant 1: It eliminates the cross-domain decoder. (2) Variant 2: It sets a uniform distribution of latent intents, with $p_(k|u_i) = 1/K$. (3) Variant 3: It omits the intra-domain contrast loss $\mathcal{L}_{orth}$, allowing the k intents to be non-orthogonal. (4) Variant 4: It removes affinity graph construction and multi-step random walks ($\tilde{R}^d = I$).

The results of DisCo and its variants are shown in Table 2. We find that: (1) Compared to Variant 1, a notable decline in performance underscores the specialized knowledge across

D	M	Model Variants				
		Variant1	Variant2	Variant3	Variant4	DisCo
Sport	H	9.81±0.24	10.64±0.19	10.62±0.37	9.56±0.27	**10.72±0.32**
	N	4.81±0.13	5.40±0.15	5.38±0.29	4.31±0.25	**5.81±0.26**
Cloth	H	11.79±0.29	12.55±0.20	12.2±0.26	11.88±0.31	**12.85±0.42**
	N	6.14±0.21	6.69±0.22	6.16±0.17	5.27±0.26	**6.92±0.32**
Game	H	9.20±0.32	9.29±0.35	9.41±0.24	8.62±0.26	**9.54±0.28**
	N	4.33±0.25	4.53±0.29	4.73±0.22	3.97±0.19	**4.87±0.27**
Video	H	11.5±0.31	12.06±0.33	12.26±0.29	12.4±0.23	**13.38±0.31**
	N	6.14±0.15	6.29±0.19	6.38±0.20	6.35±0.12	**6.86±0.14**
Music	H	14.96±0.23	15.14±0.28	13.51±0.30	13.87±0.25	**15.92±0.21**
	N	7.65±0.08	7.75±0.20	7.12±0.17	7.19±0.14	**8.56±0.13**
Movie	H	14.81±0.22	15.42±0.33	14.33±0.26	14.12±0.33	**16.33±0.28**
	N	7.49±0.17	7.87±0.19	7.31±0.20	7.33±0.24	**8.39±0.17**
Phone	H	17.27±0.29	18.21±0.35	18.22±0.30	17.70±0.34	**18.74±0.31**
	N	8.82±0.23	9.44±0.26	10.01±0.18	9.58±0.26	**10.19±0.22**
Elec	H	18.06±0.32	18.65±0.30	19.56±0.24	19.48±0.30	**21.03±0.29**
	N	9.73±0.27	10.73±0.25	10.89±0.19	10.96±0.25	**11.54±0.23**

Table 2: Ablation studies on the variants of DisCo. Dataset, Metric@10, HR, NDCG are abbreviated as D, M, H, N.

domains. This is attributed to the removal of the mapping function, which impairs the ability to learn domain-specific knowledge. (2) The comparison with Variant 2 reveals that individual users have distinct preference weights across domains, and assigning equal weight to each intent can obscure vital user preference information. (3) The comparison with Variant 3 demonstrates that maintaining orthogonality among the K intents improves prediction accuracy by capturing more diverse and richer user preference information. (4) Compared to Variant 4, the affinity graph construction and multi-step random walks leverage the inherent similarity structure of users, enhancing both intra-domain and inter-domain contrasts to filter out irrelevant collaborative information and avoid the negative transfer problem.

Parameter Analysis

We analyze the impact of three groups of hyper-parameters in DisCo: the number of latent intents, the parameter pair α and d in generating user similarity T_k (Equation 4), and the multiplier factors β and λ in the objective function (Equa-

Setting	Five Favorite Preference				
Source Domain Training	u_1: Home Entertainment u_2: Home Entertainment	Studio Specials Studio Specials	Art House & International Art House & International	Action Drama	Science Fiction Fantasy
Target Domain Test Groud Truth	u_1: Pop u_2: Pop	Metal Oldies	World Music Comedy & Spoken Word	Country Christian	Dance & Electronic Soundtracks
Target Domain Prediction by CDRIB	u_1: **Pop** u_2: **Pop**	Rock Rock	Folk Roadway	**Country** Country	Movie Soundtracks Instructional
Target Domain Prediction by DisCo	u_1: **Pop** u_2: **Pop**	Rock Jazz	New Age **Comedy & Spoken Word**	**Country** Classical	**Dance & Electronic** Symphonies

Table 3: Case Study on a user pair (u_1, u_2). Predicting results are **bold-faced** if they match the target domain test ground truth.

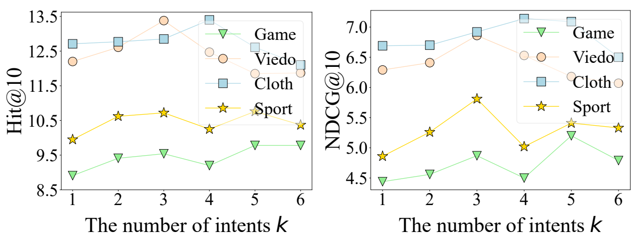

(a) Hit@10 (b) NDCG@10

Figure 3: Performance comparison w.r.t. different numbers of user intent K.

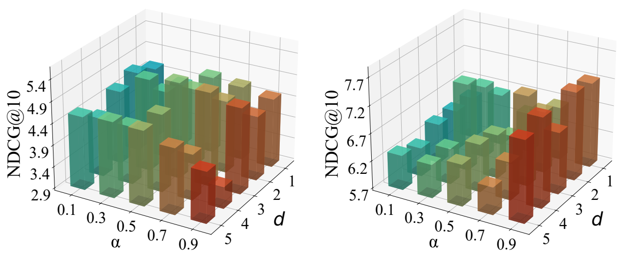

(a) Game (b) Video

Figure 4: Performance comparison w.r.t. different values of d and α for random walks and high-order user similarity.

tion 19). We find that: (1) The optimal number of latent intents is closely related to the specific domain. As shown in Figure 3, scenarios with higher interaction generally exhibit a larger number of user intents. For instance, in the domains of video and sport with dense interactions, a larger number of user intents improves performance when transferring knowledge to cold-start users in game and cloth scenarios. (2) Figure 4 demonstrates the influence of parameters α and d on generating the user similarity matrix T_k through multi-step random walks. Our results show that higher-order random walks (user similarity) generally outperform lower-order ones. Additionally, The DisCo performs well when d is set between 3 to 5 steps. (3) Figure 5 shows that an excessively high contrast loss factor λ reduces performance. And the optimal β is usually between 0.2 and 0.4.

Case Study

We conduct a case study to comfirm that DisCo effectively mitigates negative transfer. As shown in Table 3, two overlapping users (u_1,u_2) from the music-movie dataset are selected, with movies as the source domain and music as the target domain. The five most popular tags from movie interactions represent users' source domain preferences. Similarly, the five most popular tags from the top 10 ranked items generated by model depict their target domain preferences. In the target domain test ground truth, u_1 and u_2 share similar preferences in the source domain but differ in the target domain, illustrating the negative transfer issue. The results show that DisCo mitigates negative transfer, in contrast to the cold start CDR baseline, CDRIB, which fails to filter out irrelevant source-domain collaborative information.

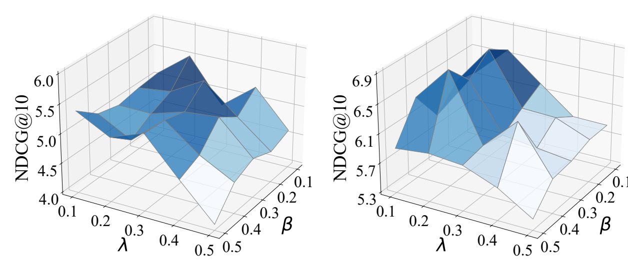

(a) Sport (b) Cloth

Figure 5: Performance comparison w.r.t. different values of λ and β for the overall objective.

Conclusion

In this paper, we propose a novel framework termed DisCo for cold-start cross-domain recommendation tasks. To address the negative transfer issue, we argue that users with similar preferences in the source domain may have different interests in the target domain, emphasizing the need to capture diverse user intents to filter irrelevant collaborative information. Specifically, a multi-channel graph encoder is employed to extract diverse user intents. Then, multi-step random walks on affinity graphs learn high-order user similarities under each intent in the source and target domains. Additionally, we treat one of the domains as the target domain and propose an intent-wise disentangled contrastive learning approach to retain user similarity relationships while explicitly identifying the rationale between the source and target domains. Experimental results on four benchmark datasets demonstrate the efficacy of our DisCo.

Acknowledgments

This paper is partially supported by the National Natural Science Foundation of China with Grant Numbers 62276002 and 62306014 as well as the Fundamental Research Funds for the Central Universities in UIBE (Grant No. 23QN02).

References

Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8): 1798–1828.

Cao, J.; Li, S.; Yu, B.; Guo, X.; Liu, T.; and Wang, B. 2023. Towards universal cross-domain recommendation. In *Proceedings of the International ACM Conference on Web Search & Data Mining*, 78–86.

Cao, J.; Lin, X.; Cong, X.; Ya, J.; Liu, T.; and Wang, B. 2022a. Disencdr: Learning disentangled representations for cross-domain recommendation. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 267–277.

Cao, J.; Sheng, J.; Cong, X.; Liu, T.; and Wang, B. 2022b. Cross-domain recommendation to cold-start users via variational information bottleneck. In *Proceedings of the IEEE Conference on Data Engineering*, 2209–2223.

Chen, X.; Duan, Y.; Houthooft, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Proceedings of the Conference on Neural Information Processing Systems*, 2172–2180.

Cheng, P.; Min, M. R.; Shen, D.; Malon, C.; Zhang, Y.; Li, Y.; and Carin, L. 2020. Improving disentangled text representation learning with information-theoretic guidance. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 7530–7541.

Guan, R.; Pang, H.; Giunchiglia, F.; Liang, Y.; and Feng, X. 2022. Cross-domain meta-learner for cold-start recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(8): 7829–7843.

Guo, X.; Li, S.; Guo, N.; Cao, J.; Liu, X.; Ma, Q.; Gan, R.; and Zhao, Y. 2023. Disentangled representations learning for multi-target cross-domain recommendation. *ACM Transactions on Information Systems*, 41(4): 1–27.

Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2016. beta-vae: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the International Conference on Learning Representations*.

Hsieh, C.-K.; Yang, L.; Cui, Y.; Lin, T.-Y.; Belongie, S.; and Estrin, D. 2017. Collaborative metric learning. In *Proceedings of the Web Conference*, 193–201.

Hu, G.; Zhang, Y.; and Yang, Q. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the International Conference on Information and Knowledge Management*, 667–676.

Ju, W.; Mao, Z.; Yi, S.; Qin, Y.; Gu, Y.; Xiao, Z.; Wang, Y.; Luo, X.; and Zhang, M. 2024a. Hypergraph-enhanced Dual Semi-supervised Graph Classification. In *Proceedings of the International Conference on Machine Learning*.

Ju, W.; Qin, Y.; Qiao, Z.; Luo, X.; Wang, Y.; Fu, Y.; and Zhang, M. 2022. Kernel-based substructure exploration for next POI recommendation. In *Proceedings of the International Conference on Data Mining*, 221–230.

Ju, W.; Wang, Y.; Qin, Y.; Mao, Z.; Xiao, Z.; Luo, J.; Yang, J.; Gu, Y.; Wang, D.; Long, Q.; et al. 2024b. Towards Graph Contrastive Learning: A Survey and Beyond. *arXiv preprint arXiv:2405.11868*.

Ju, W.; Yi, S.; Wang, Y.; Long, Q.; Luo, J.; Xiao, Z.; and Zhang, M. 2024c. A survey of data-efficient graph learning. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 8104–8113.

Ju, W.; Yi, S.; Wang, Y.; Xiao, Z.; Mao, Z.; Li, H.; Gu, Y.; Qin, Y.; Yin, N.; Wang, S.; et al. 2024d. A survey of graph neural networks in real world: Imbalance, noise, privacy and ood challenges. *arXiv preprint arXiv:2403.04468*.

Kang, S.; Hwang, J.; Lee, D.; and Yu, H. 2019. Semi-supervised learning for cross-domain recommendation to cold-start users. In *Proceedings of the International Conference on Information and Knowledge Management*, 1563–1572.

Li, B.; Yang, Q.; and Xue, X. 2009. Can movies and books collaborate? cross-domain collaborative filtering for sparsity reduction. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2052–2057.

Li, H.; Ma, W.; Sun, P.; Li, J.; Yin, C.; He, Y.; Xu, G.; Zhang, M.; and Ma, S. 2024a. Aiming at the Target: Filter Collaborative Information for Cross-Domain Recommendation. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2081–2090.

Li, H.; Wang, X.; Zhang, Z.; Yuan, Z.; Li, H.; and Zhu, W. 2021. Disentangled contrastive learning on graphs. In *Proceedings of the Conference on Neural Information Processing Systems*, 21872–21884.

Li, P.; and Tuzhilin, A. 2020. Ddtcdr: Deep dual transfer cross domain recommendation. In *Proceedings of the International ACM Conference on Web Search & Data Mining*, 331–339.

Li, X.; Sheng, J.; Cao, J.; Zhang, W.; Li, Q.; and Liu, T. 2024b. CDRNP: Cross-Domain Recommendation to Cold-Start Users via Neural Process. In *Proceedings of the International ACM Conference on Web Search & Data Mining*, 378–386.

Liang, X.; Li, D.; and Madden, A. 2020. Attributed network embedding based on mutual information estimation. In *Proceedings of the International Conference on Information and Knowledge Management*, 835–844.

Liu, J.; Sun, L.; Nie, W.; Jing, P.; and Su, Y. 2024. Graph Disentangled Contrastive Learning with Personalized Transfer for Cross-Domain Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8769–8777.

Liu, M.; Li, J.; Li, G.; and Pan, P. 2020. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the International Conference on Information and Knowledge Management*, 885–894.

Lu, Y.; Lin, Y.; Yang, M.; Peng, D.; Hu, P.; and Peng, X. 2024. Decoupled contrastive multi-view clustering with high-order random walks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 14193–14201.

Ma, J.; Cui, P.; Kuang, K.; Wang, X.; and Zhu, W. 2019a. Disentangled graph convolutional networks. In *Proceedings of the International Conference on Machine Learning*, 4212–4221.

Ma, J.; Zhou, C.; Cui, P.; Yang, H.; and Zhu, W. 2019b. Learning disentangled representations for recommendation. In *Proceedings of the Conference on Neural Information Processing Systems*, 5712–5723.

Man, T.; Shen, H.; Jin, X.; and Cheng, X. 2017. Cross-domain recommendation: An embedding and mapping approach. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2464–2470.

Qin, Y.; Wang, Y.; Sun, F.; Ju, W.; Hou, X.; Wang, Z.; Cheng, J.; Lei, J.; and Zhang, M. 2023. DisenPOI: Disentangling Sequential and Geographical Influence for Point-of-Interest Recommendation. In *Proceedings of the International ACM Conference on Web Search & Data Mining*, 508–516.

Qu, X.; Wang, Y.; Li, Z.; and Gao, J. 2024. Graph-Enhanced Prompt Learning for Personalized Review Generation. *Data Science and Engineering*, 1–16.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.

Salah, A.; Tran, T. B.; and Lauw, H. 2021. Towards source-aligned variational models for cross-domain recommendation. In *Proceedings of the ACM Conference on Recommender Systems*, 176–186.

Singh, A. P.; and Gordon, G. J. 2008. Relational learning via collective matrix factorization. In *Proceedings of the International ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 650–658.

Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019. Neural graph collaborative filtering. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 165–174.

Wang, X.; Jin, H.; Zhang, A.; He, X.; Xu, T.; and Chua, T.-S. 2020a. Disentangled graph collaborative filtering. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 1001–1010.

Wang, Y.; Li, Y.; Li, S.; Song, W.; Fan, J.; Gao, S.; Ma, L.; Cheng, B.; Cai, X.; Wang, S.; et al. 2022a. Deep graph mutual learning for cross-domain recommendation. In *International Conference on Database Systems for Advanced Applications*, 298–305.

Wang, Y.; Luo, X.; Chen, C.; Hua, X.-S.; Zhang, M.; and Ju, W. 2024. DisenSemi: Semi-supervised Graph Classification via Disentangled Representation Learning. *arXiv preprint arXiv:2407.14081*.

Wang, Y.; Qin, Y.; Han, Y.; Yin, M.; Zhou, J.; Yang, H.; and Zhang, M. 2022b. Ad-aug: Adversarial data augmentation for counterfactual recommendation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 474–490.

Wang, Y.; Qin, Y.; Sun, F.; Zhang, B.; Hou, X.; Hu, K.; Cheng, J.; Lei, J.; and Zhang, M. 2022c. DisenCTR: Dynamic graph-based disentangled representation for click-through rate prediction. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2314–2318.

Wang, Y.; Shen, J.; Song, Y.; Wang, S.; and Zhang, M. 2022d. HE-SNE: Heterogeneous event sequence-based streaming network embedding for dynamic behaviors. In *International Joint Conference on Neural Networks*, 1–8.

Wang, Y.; Song, Y.; Li, S.; Cheng, C.; Ju, W.; Zhang, M.; and Wang, S. 2022e. DisenCite: Graph-Based Disentangled Representation Learning for Context-Specific Citation Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 11449–11458.

Wang, Y.; Tang, S.; Lei, Y.; Song, W.; Wang, S.; and Zhang, M. 2020b. Disenhan: Disentangled heterogeneous graph attention network for recommendation. In *Proceedings of the International Conference on Information and Knowledge Management*, 1605–1614.

Zhang, M.; Liu, S.; and Wang, Y. 2020. STR-SA: Session-based thread recommendation for online course forum with self-attention. In *IEEE Global Engineering Education Conference*, 374–381.

Zhu, Q.; Zhu, Z.; Wang, Y.; and Yu, H. 2016. Fuzzy-information-based robustness of interconnected networks against attacks and failures. *Physica A: Statistical Mechanics and its Applications*, 458: 194–203.

Zhu, Y.; Ge, K.; Zhuang, F.; Xie, R.; Xi, D.; Zhang, X.; Lin, L.; and He, Q. 2021. Transfer-meta framework for cross-domain recommendation to cold-start users. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 1813–1817.

Zhu, Y.; Tang, Z.; Liu, Y.; Zhuang, F.; Xie, R.; Zhang, X.; Lin, L.; and He, Q. 2022. Personalized transfer of user preferences for cross-domain recommendation. In *Proceedings of the International ACM Conference on Web Search & Data Mining*, 1507–1515.