Cross-platform Recommendation Based on Prototype Alignment

Mingfeng Lv[1], Bo Li[1], Kehan Chen[1], Liukui Fan[1], Jixiang Lei[1]

[1]AI Class, Xiamen University
Xiamen, China
{23020241154429, 23020241154349, 23020241154340, 23020241154388, 36920241153225}@stu.xmu.edu.cn

Abstract

Cross-platform recommendation (CPR) has attracted increasing attention as a solution to the problem of data sparsity by leveraging information from different platforms with similar service or items. The current CPR methods concentrate on the transfer of user and item representations for the purpose of augmenting cross-platform information. However, these methods face two main challenges: the divergence in the item space across platforms and the divergence in the distribution of user preferences across platforms. To address these challenges simultaneously, we propose a CPR framework based on prototype alignment (CPPA). Firstly, from item perspective, our framework constructs a unified semantic space for items based on the Large Language Model (LLM) and transfers information of cross-platform items through prototype alignment. Second, from user perspective, our framework builds and aligns user preference prototypes across platforms. Notably, CPPA is a general framework that can easily extend existing single-platform collaborative filtering models to CPR.

Introduction

Recommendation Systems (RSs) are crucial for mitigating information overload in areas like e-commerce (Liu et al. 2023; Wang et al. 2023), advertising (Bai et al. 2022; Fan, Si, and Zhang 2023), and social networking (Quan et al. 2023; Zhang et al. 2023b). However, data sparsity from limited feedback remains a significant challenge. Cross-Platform Recommendation (CPR) and Cross-Domain Recommendation (CDR) have emerged as effective solutions by leveraging multi-platform or cross-domain information (Zang et al. 2023).

Existing CDR methods often use overlapping user or item information for knowledge transfer (Xu et al. 2023; Zhao et al. 2023a), but privacy concerns make such approaches impractical in CPR, where data does not overlap. CPR techniques typically rely on textual content transfer (Li et al. 2023; Liu et al. 2022a; Yu et al. 2020; Zhang et al. 2019) or aligning global user preferences (Li et al. 2022a; Zhang et al. 2023a; Zhao et al. 2023b). Despite these advancements, current methods struggle with cross-platform divergence and bias, leaving critical issues underexplored.

Figure 1: Item ID and textual content divergence on Amazon and IMDB.

Challenge 1: Item Space Divergence. Cross-platform item divergence arises from two main factors: (1) **Item ID Divergence.** Item IDs differ across platforms, hindering direct feedback transfer. For instance, the movie *Titanic* has distinct IDs on Amazon and IMDB (Figure 1), making ID-based methods ineffective. (2) **Item Content Divergence.** Existing methods rely on item textual content for cross-platform transfer (Choi et al. 2022; Liu et al. 2022a; Yu et al. 2020; Li et al. 2023; Zhang et al. 2019), but descriptions and reviews vary greatly. Amazon provides detailed plot summaries and user reviews focused on viewing experience, whereas IMDB offers brief summaries and emotionally-driven reviews. This semantic divergence leads to inconsistent representations and negative transfer effects.

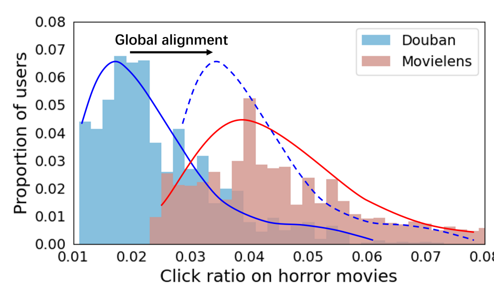

Figure 2: Divergence in user preference distributions across platforms.

Challenge 2: Divergence in User Preference Distributions. Existing methods (Zhao et al. 2023b; Zhang et al. 2023a; Li et al. 2022a) aim to transfer global user preferences across platforms. However, as shown in Figure 2, these approaches suffer from misalignment due to divergent user preference distributions. Specifically, platforms have user preference groups of varying sizes but corresponding segments. For example, the range $[0.3, 0.5]$ on the x-axis represents the same preference segment on Movielens and Douban. While alignment should occur here, global methods shift the entire distribution—moving it right on Douban and left on Movielens—exacerbating preference gaps and hindering effective transfer across platforms.

Relative Work

Cross-Domain Recommendation

In the era of big data, despite the massive volume of data, domain- or user-specific data is often limited, leading to issues such as *data sparsity* and the *cold start problem*. Data sparsity reduces recommendation accuracy, while the cold start problem arises from insufficient historical data for new users or items. To address these challenges, cross-domain recommendation techniques have been proposed (Zang et al. 2023), leveraging relationships and similarities across domains to enhance performance in the target domain. For example, HCCDR (Xu et al. 2023) clusters overlapping users to learn personalized biases, assuming similar interests across domains. NATR (Gao et al. 2022) transfers item information by training a unified neural network to learn inter-domain item relationships. M3Rec (Cao et al. 2022) integrates overlapping item information and explores both direct and indirect cross-domain connections. While effective, these methods heavily rely on overlapping information, such as shared user or item IDs, which limits their applicability in cross-platform scenarios where such overlaps are often absent. For instance, users may use different IDs across platforms, or items may have distinct identifiers, making cross-domain techniques difficult to extend to cross-platform recommendation tasks.

Cross-Platform Recommendation

To overcome the reliance on overlapping information in cross-domain recommendation, cross-platform recommendation (CPR) has become a key research focus, utilizing diverse techniques for cross-platform information transfer, primarily through representation transfer. *Global user interaction transfer* captures shared characteristics across platforms. For example, ALCDR (Zhao et al. 2023b) aggregates user representations via similarity matrices, GWCDR (Li et al. 2022a) aligns cross-domain features to enhance target recommendations, and UniCDR (Cao et al. 2023) employs contrastive learning for improved shared feature representation. When interaction data is unavailable, *textual content transfer* acts as a bridge. SCT (Zhang et al. 2019) leverages semantic tags for alignment, while TDAR (Yu et al. 2020) and CFAA (Liu et al. 2022a) utilize textual similarities to link latent spaces across platforms. These methods eliminate the dependence on overlapping data, offering innovative solutions for accurate and personalized recommendations in cross-platform scenarios.

Methodology

Problem Definition

A cross-platform recommendation (CPR) framework is proposed, involving two platforms: the source platform $D^s = U^s, V^s, R^s, T^s$ and the target platform $D^t = U^t, V^t, R^t, T^t$. Here, U^s and U^t are user sets, V^s and V^t are item sets, and T^s and T^t are text sets for the source and target platforms, respectively. R^s and R^t denote interaction matrices, where $R_{u,v} = 1$ indicates an interaction between user u and item v, and $R_{u,v} = 0$ indicates no interaction.

Overall Framework

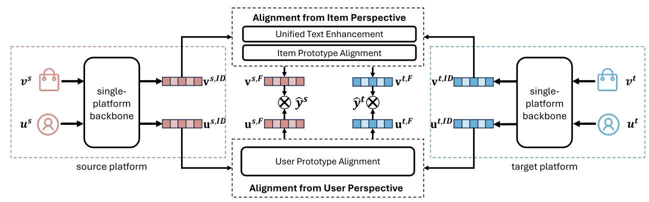

Figure 3: The overall framework of CPPA.

To address challenges from both item and user perspectives and improve backbone network performance on a single platform, this paper proposes CPPA, a cross-platform recommendation framework based on prototype alignment. As shown in Figure 3, the framework consists of two key components: item alignment and user alignment. **(1) Item Alignment.** This component includes a unified text enhancement module and an item prototype alignment module. First, item titles are generated using a large language model, creating unified text embeddings that enhance user ID embeddings. Item prototypes are then constructed and aligned to link cross-platform items. **(2) User Alignment.** This component features a user prototype alignment module, enabling the transfer of cross-platform user preferences without relying on overlapping user data.

Unified Text Enhancement

To address item text divergences across platforms, directly using original text for cross-platform embeddings often leads to negative transfer due to inconsistent embedding spaces. To mitigate this, we use large language models (LLMs) to generate unified item text, ensuring embeddings lie in a shared semantic space for high-quality alignment. The implementation is shown in Figure 4.

First, the item title is input to the LLM, which generates semantically rich text using its knowledge base. Task-specific prompts are designed to ensure the text captures essential item details. For example, for movies, the LLM generates content on plot, actors, directors, strengths, and weaknesses. A one-shot prompting strategy is used to enforce consistent formatting, reducing cross-platform divergences.

Next, embeddings are extracted from the unified text using the CLS token $\mathbf{v}^B$ from the pre-trained BERT model (Devlin et al. 2019), representing the semantic content. A shared mapping layer $Proj$ refines these embeddings to incorporate recommendation-related information, producing the final embedding $\mathbf{v}^T$. This embedding enhances the item ID embedding $\mathbf{v}^{ID}$ from the backbone network. The final enhanced item embedding $\mathbf{v}^E$ is computed as:

$$\mathbf{v}^E = (1-\alpha)\mathbf{v}^T + \alpha\mathbf{v}^{ID}, \tag{1}$$

where α is a hyperparameter balancing the two embeddings.

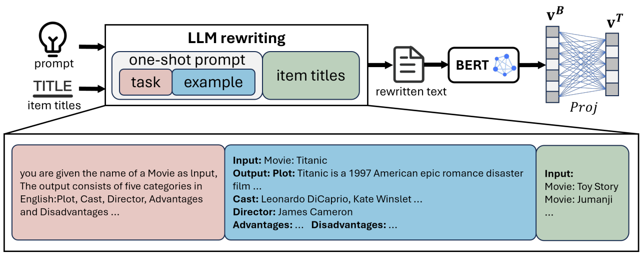

Figure 4: The unified text enhancement module.

Item Prototype Alignment

Due to differing item ID settings across platforms, direct cross-platform connections via item IDs are infeasible. To address this, the proposed model introduces item prototype alignment, which aligns the clustering centers of items across platforms. By summarizing item features, prototypes enable the connection of similar items without requiring overlapping cross-platform item IDs.

As shown in Figure 5, the module consists of three components: prototype initialization, intra-platform alignment, and inter-platform alignment. To ensure prototype stability, initialization occurs only when the main network achieves a relatively stable state, typically when validation performance begins to decline.

(1) Prototype Initialization. Prototype matrices $\mathbf{C}^s$ and $\mathbf{C}^t$ for platforms D^s and D^t are initialized to record embeddings of item clusters (i.e., prototypes). These matrices are defined as $\mathbf{C} = \{\mathbf{c}_1, \mathbf{c}_2, ..., \mathbf{c}_P\} \in \mathbb{R}^{L\times P}$, where P is the number of prototypes, and $\mathbf{c}_i \in \mathbb{R}^L$ is the learnable embedding for the i-th prototype. To avoid unmatched prototypes from random initialization, clustering centers are used as initial values. Item data from both platforms are unified, clustered, and the cluster centers serve as prototype initializations.

(2) Intra-platform Alignment. For each item embedding $\mathbf{v}^E$ on a platform, the corresponding prototype $\mathbf{c}$ is updated to better reflect shared item features. The intra-platform alignment loss minimizes the distance between item embeddings and their prototypes:

$$L^{item}_{intra} = -\sum_{v\in V}(||sg(\mathbf{c}) - \mathbf{v}^E||^2 + ||\mathbf{c} - sg(\mathbf{v}^E)||^2), \tag{2}$$

where sg is the stop gradient operator.

(3) Inter-platform Alignment. After obtaining prototypes for both platforms, inter-platform alignment transfers knowledge of cross-platform item prototypes. The alignment loss is defined as:

$$\begin{aligned} L^{item}_{inter} = &-\sum_{i=1}^{P} \log \frac{\exp(c^s_i \cdot sg(c^t_i)/\tau)}{\sum_{i'=1}^{P}\exp(c^s_{i'} \cdot sg(c^t_i)/\tau)} \\ &-\sum_{i=1}^{P} \log \frac{\exp(c^t_i \cdot sg(c^s_i)/\tau)}{\sum_{i'=1}^{P}\exp(c^t_{i'} \cdot sg(c^s_i)/\tau)}, \end{aligned} \tag{3}$$

where τ is a temperature parameter.

The total prototype alignment loss is:

$$L^{item} = L^{item}_{intra} + L^{item}_{inter}. \tag{4}$$

To further enhance recommendation accuracy, prototype embeddings are integrated into item embeddings:

$$\mathbf{v}^F = (1-\beta)\mathbf{v}^E + \beta\mathbf{c}, \tag{5}$$

where $\mathbf{v}^F$ is the final item embedding, and β is a hyperparameter.

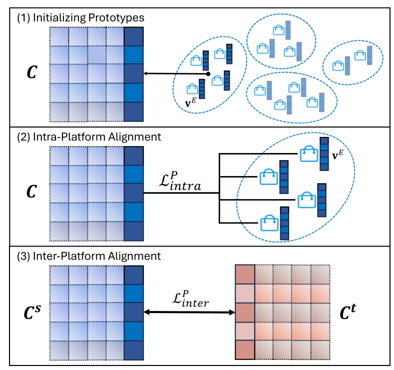

Figure 5: The item prototype alignment module.

User Prototype Alignment

Existing global alignment methods (Zhao et al. 2023b; Zhang et al. 2023a; Li et al. 2022a) assume transferable global user preference distributions, which is unrealistic as preferences often vary significantly across platforms, increasing the gap between overlapping preferences.

To address this, the proposed user prototype alignment method extracts and aligns core user preference features, independent of distribution differences. The alignment process mirrors that of item prototypes. The user prototype alignment loss, L^{user}, and final embedding, $\mathbf{u}^F$, are derived accordingly.

Optimization

The original single-platform recommendation system typically employs the BPR loss L^{BPR} as the supervised loss, which is also adopted in this work. The BPR loss is defined as:

$$L^{BPR} = -\sum_{(u,v_p,v_n)\in B} \ln\sigma(\mathbf{u}^F\mathbf{v}^F_p - \mathbf{u}^F\mathbf{v}^F_n), \tag{6}$$

where $(u, v_p, v_n) \in B$ represents a triplet, with v_p as a positive sample (indicating interaction with u) and v_n as a negative sample. σ denotes the sigmoid function.

Table 1: Overall performance of cross-platform recommendation. The best improvements are highlighted in bold. %Improv. represents the percentage of improvement of CPPA to the best result in corresponding competitors.

Method		*MovieLens-1m→Douban Movies*				*Douban Movies→MovieLens-1m*			
		HR@10	NDCG@10	Recall@10	Precision@10	HR@10	NDCG@10	Recall@10	Precision@10
CPR models	SCDGN	0.0297	0.0394	0.0394	0.0288	0.1149	0.1064	0.1442	0.0499
	UniCDR	0.0317	0.0150	0.0471	0.0307	0.1024	0.0489	0.1450	0.0445
CPR plug-and-play frameworks	MF	0.0312	0.0402	0.0403	0.0303	0.1340	0.1253	0.1673	0.0582
	+ SRTrans	0.0194	0.0263	0.0275	0.0188	0.1145	0.1014	0.1375	0.0498
	+ GWCDR	0.0324	0.0409	0.0372	0.0314	0.1315	0.1181	0.1561	0.0571
	+ CPPA	**0.0374**	**0.0508**	**0.0485**	**0.0362**	**0.1557**	**0.1453**	**0.1960**	**0.0677**
	%Improv.	15.14%	24.26%	20.39%	15.14%	16.17%	15.95%	17.15%	16.17%
	LightGCN	0.0463	0.0644	0.0621	0.0448	0.1729	0.1630	0.2113	0.0751
	+ SRTrans	0.0481	0.0655	0.0602	0.0466	0.1646	0.1528	0.2009	0.0716
	+ GWCDR	0.0452	0.0610	0.0568	0.0437	0.1730	0.1621	0.2125	0.0752
	+ CPPA	**0.0510**	**0.0699**	**0.0670**	**0.0494**	**0.1801**	**0.1694**	**0.2233**	**0.0783**
	%Improv.	6.01%	6.71%	7.76%	6.01%	4.10%	3.93%	5.07%	4.10%
Method		*MovieTwettings→Douban Movies*				*Douban Movies→MovieTwettings*			
		HR@10	NDCG@10	Recall@10	Precision@10	HR@10	NDCG@10	Recall@10	Precision@10
CPR models	SCDGN	0.0302	0.0382	0.0398	0.0293	0.1130	0.0773	0.1328	0.0213
	UniCDR	0.0310	0.0147	0.0477	0.0300	0.1166	0.0568	0.1399	0.0220
CPR plug-and-play frameworks	MF	0.0312	0.0402	0.0403	0.0303	0.1318	0.0884	0.1493	0.0249
	+ SRTrans	0.0196	0.0251	0.0261	0.0190	0.0657	0.0428	0.0705	0.0124
	+ GWCDR	0.0316	0.0399	0.0365	0.0306	0.1044	0.0673	0.1138	0.0197
	+ CPPA	**0.0365**	**0.0493**	**0.0473**	**0.0353**	**0.1395**	**0.0955**	**0.1585**	**0.0263**
	%Improv.	15.40%	22.66%	17.47%	15.40%	5.79%	8.03%	6.14%	5.79%
	LightGCN	0.0463	0.0644	0.0621	0.0448	0.1487	0.1001	0.1655	0.0281
	+ SRTrans	0.0483	0.0650	0.0613	0.0467	0.1385	0.0933	0.1572	0.0261
	+ GWCDR	0.0455	0.0614	0.0574	0.0441	0.1492	0.1025	0.1688	0.0281
	+ CPPA	**0.0519**	**0.0693**	**0.0668**	**0.0502**	**0.1603**	**0.1076**	**0.1797**	**0.0302**
	%Improv.	7.53%	6.57%	7.58%	7.53%	7.46%	5.00%	6.47%	7.46%

Table 2: Performance of different components of CPPA

Method		*MovieLens-1m→Douban Movies*				*Douban Movies→MovieLens-1m*			
		HR@10	NDCG@10	Recall@10	Precision@10	HR@10	NDCG@10	Recall@10	Precision@10
MF	Vanilla	0.0312	0.0402	0.0403	0.0303	0.1340	0.1253	0.1673	0.0582
	+ UTE	0.0357	0.0467	0.0473	0.0345	0.1460	0.1356	0.1822	0.0635
	+ UTE & IPA	0.0359	0.0484	0.0476	0.0348	0.1485	0.1379	0.1853	0.0645
	+ UTE & IPA & UPA	**0.0374**	**0.0508**	**0.0485**	**0.0362**	**0.1557**	**0.1453**	**0.1960**	**0.0677**
LightGCN	Vanilla	0.0463	0.0644	0.0621	0.0448	0.1729	0.1630	0.2113	0.0751
	+ UTE	0.0483	0.0655	0.0636	0.0468	0.1768	0.1677	0.2182	0.0768
	+ UTE & IPA	0.0492	0.0669	0.0646	0.0476	0.1782	0.1680	0.2206	0.0774
	+ UTE & IPA & UPA	**0.0510**	**0.0699**	**0.0670**	**0.0494**	**0.1801**	**0.1694**	**0.2233**	**0.0783**

The overall loss function integrates the BPR loss L^{BPR}, item prototype alignment loss L^{item}, and user prototype alignment loss L^{user}:

$$L = L^{BPR} + \lambda(L^{item} + L^{user}), \tag{7}$$

where λ balances the loss components. This optimization step enhances the backbone network by incorporating cross-platform knowledge.

Experiment

In this section, the proposed CPPA framework is evaluated on two real-world datasets using two different backbone networks. The main research questions are as follows:

Question 1: Can this framework improve the performance of the backbone network in cross-platform recommendation?

Question 2: What is the performance of each module in this framework?

Question 3: How does the framework perform on sparser datasets?

Experimental Setup

Datasets: We evaluate the proposed framework on two dataset pairs: *MovieLens-1m* and *Douban Movies*, *MovieTweetings* and *Douban Movies*. These pairs represent typical cross-platform recommendation scenarios with non-overlapping users and items, but similar services or item types.

Following prior work (Liu et al. 2022b; Li et al. 2022a), ratings are mapped to binary values (0 or 1). Users with at least 5 interactions and items with at least 10 interactions are retained. For denser movie datasets, ratings of 5 (scale of 5) are mapped to 1, while others are mapped to 0. This ensures the interaction matrix captures meaningful preferences while avoiding excessive sparsity. Each dataset is split into 80% training, 10% validation, and 10% testing.

Evaluation Metrics: The evaluation employs widely used metrics: HR@k (Hit Rate@k), NDCG@k, Recall@k, and Precision@k, with $k = 10$.

Implementation Details: The framework is implemented in PyTorch, with both user and item embeddings set to 64 dimensions. The Adam optimizer is used, and the backbone network architecture remains consistent across frameworks.

Baselines

The proposed framework is a general cross-platform recommendation system designed for single-platform collaborative filtering backbones. For comparison, we evaluate the framework using two backbone networks, MF (Koren, Bell, and Volinsky 2009) and LightGCN (He et al. 2020), and compare it with state-of-the-art cross-platform framework. A brief description of these baselines is provided below.

(1) Single-Platform Backbone Networks: MF (Koren, Bell, and Volinsky 2009) is a classical matrix factorization model for single-platform recommendations. LightGCN (He et al. 2020) simplifies graph convolutional networks to improve single-platform recommendation performance.

(2) Cross-Platform General Frameworks: GWCDR (Li et al. 2022a) aligns representation distributions using Gromov-Wasserstein distance, enabling knowledge transfer across platforms. SRTrans (Li et al. 2023) extracts relational knowledge between platforms via semantic clustering of source and target items.

(3) Cross-Platform Models: SCDGN (Li et al. 2022b) constructs a cross-platform bipartite graph using semantic item clustering and user interactions. UniCDR (Cao et al. 2023) employs masking and aggregation mechanisms to transfer shared information across platforms.

Main Results

This paper evaluates CPPA on two real-world dataset pairs, with results shown in Table 1. Key observations are as follows.

(1) **CPPA significantly enhances MF and LightGCN.** For MF, improvements by CPPA are 12.95% (HR@10), 15.67% (NDCG@10), 13.93% (Recall@10), and 12.95% (Precision@10). For LightGCN, improvements are 9.02%, 8.68%, 9.46%, and 9.02%, respectively. Despite LightGCN's strong performance, CPPA achieves notable gains by transferring cross-platform information.

(2) **CPPA performs well across all datasets.** On *MovieLens-1m* and *MovieTweetings*, improvements are 12.99% (HR@10) and 14.90% (NDCG@10). On *MovieTweetings* and *MovieLens-1m*, improvements are 11.11% and 12.60%. Larger gains on these datasets highlight CPPA's ability to transfer cross-platform knowledge to diverse items.

(3) **CPPA outperforms all baselines.** Compared to SCDGN, UniCDR, SRTrans, and GWCDR, CPPA shows consistent improvements, with average gains of 9.15% (HR@10) and 10.77% (NDCG@10). Its ability to mitigate negative transfer from item and preference differences ensures superior performance.

In summary, CPPA mitigates negative transfer via prototype alignment, consistently improving recommendation performance across datasets.

Ablation Study

To assess the effectiveness of each module in CPPA, we conduct an ablation study on *Douban Movies* and *MovieLens-1m*.

This study evaluates the contributions of three key modules: Unified Text Enhancement **(UTE)**, Item Prototype Alignment **(IPA)**, and User Prototype Alignment **(UPA)**. Performance is evaluated incrementally as modules are added to the backbone model. Results are shown in Table 2. Key observations are as follows.

(1) **All modules contribute to performance improvement.** Sequential addition of UTE, IPA, and UPA leads to average improvements of 8.25%, 1.41%, and 3.45%, respectively, highlighting the positive impact of each module in transferring cross-platform knowledge.

(2) **UTE provides the most significant gains.** UTE, with an 8.25% average improvement, effectively addresses cross-platform text divergence and enriches item semantics, enhancing item attribute and relationship understanding. Other modules further improve performance, demonstrating CPPA's ability to leverage cross-platform information fully.

Performance on Sparse Datasets

Figure 6 presents the results. Across all sparsity levels, LightGCN+CPPA consistently outperforms LightGCN, demonstrating its ability to leverage cross-platform information to alleviate data sparsity. Notably, as sparsity increases, the relative improvement of LightGCN+CPPA becomes more pronounced, highlighting its effectiveness in capturing cross-platform user and item features. This capability provides richer and more accurate information, significantly enhancing recommendation performance in sparse conditions.

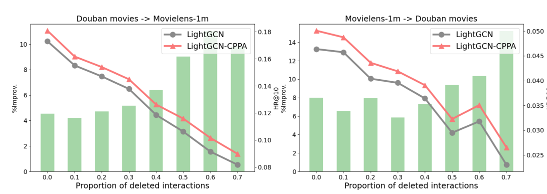

Figure 6: Performance of LightGCN and LightGCN+CPPA on datasets with different sparsity levels. The green bars indicate the improvement of LightGCN+CPPA.

Conclusion

In this paper, we identify two under-explored challenges in the cross-platform recommendation (CPR): the divergence in the item space and the divergence in the distribution of user preference. We propose a CPR framework based on Prototype Alignment (CPPA) to tackle these challenges from the perspectives of both items and users. Moreover, CPPA is a plug-and-play framework that can be applied to single-platform collaborative filtering backbones and simultaneously enhance performance on both platforms. Extensive experiments on two pairs of real datasets show that CPPA provides a stable and significant improvement, outperforming SOTA CPR frameworks and models.

References

Bai, X.; Duan, L.; Tang, R.; Batra, G.; and Agrawal, R. 2022. Improving Text-based Similar Product Recommendation for Dynamic Product Advertising at Yahoo. In *CIKM*, 2883–2892.

Cao, J.; Cong, X.; Liu, T.; and Wang, B. 2022. Item Similarity Mining for Multi-Market Recommendation. In *SIGIR*, 2249–2254.

Cao, J.; Li, S.; Yu, B.; Guo, X.; Liu, T.; and Wang, B. 2023. Towards Universal Cross-Domain Recommendation. In *WSDM*, 78–86.

Choi, Y.; Choi, J.; Ko, T.; Byun, H.; and Kim, C. 2022. Review-Based Domain Disentanglement without Duplicate Users or Contexts for Cross-Domain Recommendation. In *CIKM*, 293–303.

Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*, 4171–4186.

Fan, Y.; Si, N.; and Zhang, K. 2023. Calibration Matters: Tackling Maximization Bias in Large-scale Advertising Recommendation Systems. In *ICLR*.

Gao, C.; Li, Y.; Feng, F.; Chen, X.; Zhao, K.; He, X.; and Jin, D. 2022. Cross-domain Recommendation with Bridge-Item Embeddings. *ACM Trans. Knowl. Discov. Data*, 16(1): 2:1–2:23.

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.

Koren, Y.; Bell, R. M.; and Volinsky, C. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8): 30–37.

Li, X.; Qiu, Z.; Zhao, X.; Wang, Z.; Zhang, Y.; Xing, C.; and Wu, X. 2022a. Gromov-Wasserstein Guided Representation Learning for Cross-Domain Recommendation. In *CIKM*, 1199–1208. ACM.

Li, Z.; Amagata, D.; Zhang, Y.; Hara, T.; Haruta, S.; Yonekawa, K.; and Kurokawa, M. 2022b. Debiasing Graph Transfer Learning via Item Semantic Clustering for Cross-Domain Recommendations. In *IEEE Big Data*, 762–769.

Li, Z.; Amagata, D.; Zhang, Y.; Hara, T.; Haruta, S.; Yonekawa, K.; and Kurokawa, M. 2023. Semantic Relation Transfer for Non-Overlapped Cross-Domain Recommendations. In Kashima, H.; Ide, T.; and Peng, W.-C., eds., *Advances in Knowledge Discovery and Data Mining*, 271–283. Cham. ISBN 978-3-031-33380-4.

Liu, W.; Zheng, X.; Hu, M.; and Chen, C. 2022a. Collaborative Filtering with Attribution Alignment for Review-based Non-overlapped Cross Domain Recommendation. In *WWW*, 1181–1190.

Liu, W.; Zheng, X.; Hu, M.; and Chen, C. 2022b. Collaborative Filtering with Attribution Alignment for Review-based Non-overlapped Cross Domain Recommendation. In *WWW*, 1181–1190.

Liu, Y.; Zhang, W.; Dong, B.; Fan, Y.; Wang, H.; Feng, F.; Chen, Y.; Zhuang, Z.; Cui, H.; Li, Y.; and Che, W. 2023. U-NEED: A Fine-grained Dataset for User Needs-Centric E-commerce Conversational Recommendation. In *SIGIR*, 2723–2732.

Quan, Y.; Ding, J.; Gao, C.; Yi, L.; Jin, D.; and Li, Y. 2023. Robust Preference-Guided Denoising for Graph based Social Recommendation. In *WWW*, 1097–1108.

Wang, Z.; Zou, Y.; Dai, A.; Hou, L.; Qiao, N.; Zou, L.; Ma, M.; Ding, Z.; and Xu, S. 2023. An Industrial Framework for Personalized Serendipitous Recommendation in E-commerce. In *RecSys*, 1015–1018.

Xu, J.; Wang, X.; Zhang, H.; and Lv, P. 2023. Heterogeneous and clustering-enhanced personalized preference transfer for cross-domain recommendation. *Inf. Fusion*, 99: 101892.

Yu, W.; Lin, X.; Ge, J.; Ou, W.; and Qin, Z. 2020. Semi-supervised Collaborative Filtering by Text-enhanced Domain Adaptation. In *KDD*, 2136–2144.

Zang, T.; Zhu, Y.; Liu, H.; Zhang, R.; and Yu, J. 2023. A Survey on Cross-domain Recommendation: Taxonomies, Methods, and Future Directions. *ACM Trans. Inf. Syst.*, 41(2): 42:1–42:39.

Zhang, Q.; Hao, P.; Lu, J.; and Zhang, G. 2019. Cross-Domain Recommendation with Semantic Correlation in Tagging Systems. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8.

Zhang, Q.; Liao, W.; Zhang, G.; Yuan, B.; and Lu, J. 2023a. A Deep Dual Adversarial Network for Cross-Domain Recommendation. *IEEE Trans. Knowl. Data Eng.*, 35(4): 3266–3278.

Zhang, X.; Xu, S.; Lin, W.; and Wang, S. 2023b. Constrained Social Community Recommendation. In *KDD*, 5586–5596.

Zhao, C.; Zhao, H.; He, M.; Zhang, J.; and Fan, J. 2023a. Cross-domain recommendation via user interest alignment. In *WWW*, 887–896.

Zhao, Y.; Li, C.; Peng, J.; Fang, X.; Huang, F.; Wang, S.; Xie, X.; and Gong, J. 2023b. Beyond the Overlapping Users: Cross-Domain Recommendation via Adaptive Anchor Link Learning. In *SIGIR*, 1488–1497.