

ABXI: Invariant Interest Adaptation for Task-Guided Cross-Domain Sequential Recommendation

Anonymous Author(s)
Submission Id: 227

ABSTRACT

Cross-Domain Sequential Recommendation (CDSR) has recently gained attention for countering data sparsity by transferring knowledge across domains. A common approach merges domain-specific sequences into cross-domain sequences, serving as bridges that enable mutual enhancement between domains. One key challenge is to correctly extract the effective shared knowledge among these sequences and appropriately transfer it. Most existing works directly transfer unfiltered cross-domain knowledge rather than extracting domain-invariant components and adaptively integrating them into domain-specific modelings. Another challenge lies in aligning the domain-specific and cross-domain sequences. Existing methods align these sequences based on timestamps, but this approach can cause prediction mismatches when the current tokens and their targets belong to different domains. In such cases, the domain-specific knowledge carried by the current tokens may degrade performance. To address these challenges, we propose the A-B-Cross-to-Invariant Learning Recommender (ABXI). Specifically, leveraging LoRA's effectiveness for efficient adaptation as supported by numerous studies, our model incorporates two types of LoRAs to facilitate the adaptation process. First, all sequences are processed through a shared encoder that employs a domain LoRA for each sequence, thereby preserving unique domain characteristics. Next, we introduce an invariant projector that extracts domain-invariant interests from cross-domain representations, utilizing an invariant LoRA as well to adapt these interests into recommendations in each specific domain. Besides, to avoid prediction mismatches, all domain-specific sequences are re-aligned to match the domains of the cross-domain ground truths. Experimental results on three datasets demonstrate that our approach achieves better results than other CDSR counterparts, with an average improvement of 17.30% in HR@10 and 18.65% in NDCG@10. The codes are available in <https://anonymous.4open.science/status/ABXI-WWW25-1D04>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender Systems, Cross-Domain Sequential Recommendation, Low-Rank Adaptation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2024 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Anonymous Author(s). 2024. ABXI: Invariant Interest Adaptation for Task-Guided Cross-Domain Sequential Recommendation. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In the era of information explosion, the Internet is flooded with massive amounts of content, yet users are exposed to only a small fraction of it. Such data sparsity remains a persistent challenge in modern recommender systems. Cross-Domain Sequential Recommendation (CDSR) has recently emerged as a promising approach to alleviate this sparsity issue by transferring knowledge across different domains to enrich user profiles [4, 5, 9, 11, 30, 41, 53, 55, 60].

A common strategy in CDSR involves merging domain-specific sequences into cross-domain sequences that serve as bridges, enabling mutual enhancement between domains [4, 9, 30, 53, 55]. Figure 1 illustrates an example of a user's domain-specific and cross-domain interaction sequences. In the book domain, the user's interests encompass science fiction and romantic novels, while in the movie domain, the user prefers science fiction and comedy films. From the perspective of cross-domain sequences, the user's interest in science fiction can be leveraged in both the book and movie domains to create more comprehensive user profiles. On the contrary, the specific interests in romantic books and comedy movies should not be indiscriminately shared between domains. However, most existing CDSR approaches mix up the concepts of cross-domain and domain-invariant interests by directly transferring unfiltered cross-domain knowledge into domain-specific modeling. This practice can introduce domain-specific information from one domain to interfering with another, adversely affecting recommendation performance. Therefore, extracting domain-invariant knowledge from cross-domain sequences is essential to facilitate effective sharing across specific domains.

Furthermore, another challenge lies in aligning domain-specific and cross-domain sequences when making recommendations within each domain. Current self-attention-based methods typically align cross-domain and domain-specific sequences based on timestamps [9, 30, 55], as depicted in Figure 2b. Although this approach is intuitive and facilitates the enhancement of cross-domain sequential features with domain-specific features through direct token-wise addition, it has inherent limitations. Generally, cross-domain training sequences comprise a mixture of items from different domains. If the current token and its ground truth token belong to different domains, the domain-specific information encoded in the current token may negatively impact the prediction of the ground truth token. For instance, as illustrated in Figure 2b, consider selecting token A4 as the current token to predict token B4. The timestamp-guided alignment enables the model to incorporate encoded cross-domain interests along seq_x as well as encoded domain-specific interests

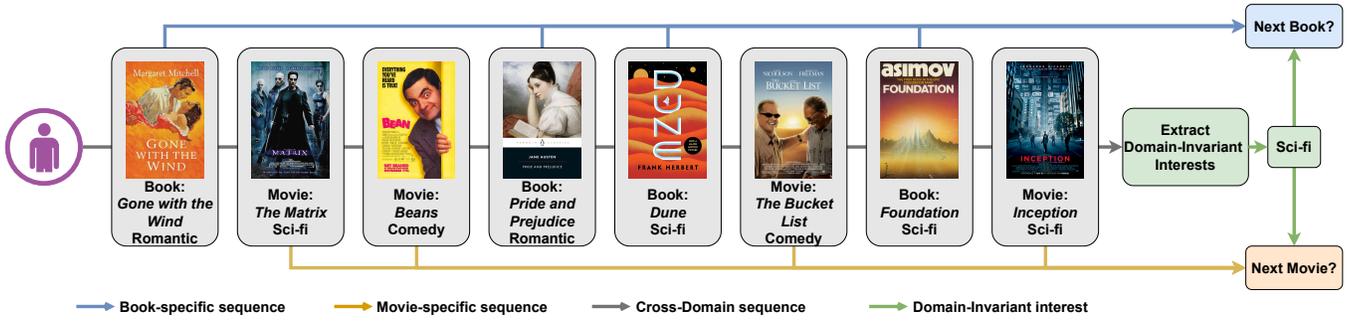


Figure 1: Our proposal on generating recommendations by integrating domain-specific interests with domain-invariant interests extracted from the cross-domain sequence.

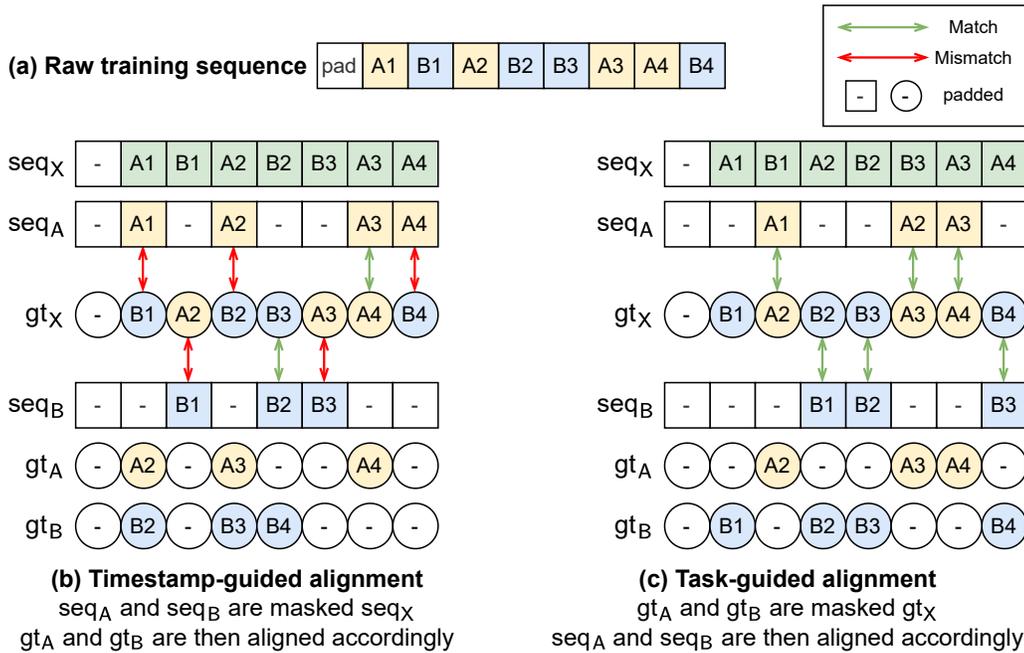


Figure 2: Illustration of the sequence splits under different alignments, where gt denotes the ground truth. (a) illustrates the input training sequence. (b) and (c) demonstrate the split outcomes of timestamp-guided and task-guided alignment, respectively.

along seq_A. However, since the target token B4 originates from domain B, it does not correspond with the A-specific knowledge, potentially degrading the model’s performance. We refer to this issue as the prediction mismatch throughout the remainder of this paper.

To tackle the prediction mismatch issue and address the challenges of exploiting domain-invariant interest, we propose the A-B-Cross-to-Invariant Learning Recommender (ABXI). Specifically, we first realign all domain-specific sequences according to the domains of the ground truths to prevent prediction mismatches with cross-domain sequences. We then employ a shared self-attention encoder as the sequence model to encode all sequences into sequential representations. This shared encoder deploys one domain LoRA (dLoRA)

for each sequence, which can efficiently switch modes to encode every cross-domain and domain-specific sequence. Additionally, we instantiate an invariant projector to extract the domain-invariant interests from cross-domain representations. This projector has integrated one invariant LoRA (iLoRA) for recommendations in each specific domain to conduct efficient adaptation. While LoRAs are typically used for fine-tuning, we extend their application to single-stage training by concurrently training LoRA modules with all other components in ABXI. Having introduced these designs, ABXI renovates both the pipelines of obtaining cross-domain and domain-specific interests.

To thoroughly evaluate ABXI with state-of-the-art CDSR methods, we conduct extensive experiments on three publicly available

datasets. Experimental results show that ABXI outperforms all baselines by a significant margin, achieving notable improvements of 17.30% on HR@10 and 18.65% on NDCG@10. Ablation studies and sensitivity analyses further demonstrate the effectiveness of our proposed designs.

To conclude, our contributions can be summarized as follows:

- We identify the prediction mismatch problem within previous sequence-model-based CDSR works, and introduce a task-guided alignment to solve this problem.
- We introduce two types of LoRA: dLoRAs switch the mode of the encoder to handle encoding each sequence; iLoRAs adaptively integrate domain-invariant interests into recommendations in each specific domain.
- Extensive experiments are provided to demonstrate the effectiveness of ABXI. Results show that ABXI outperforms all baselines including state-of-the-art CDSR counterparts.

The rest of this paper is organized as follows: Section 2 provides an overview of related work. In Section 3, we formalize the CDSR problem we aim to solve and introduce our proposed ABXI. Section 4 evaluates ABXI through extensive experiments; Section 5 presents the conclusion.

2 RELATED WORK

2.1 Cross-Domain Recommendation

Cross-Domain Recommendation (CDR) leverages transfer learning techniques to mitigate data sparsity. Common methods include domain alignment, which aligns users' or items' representations across different domains [36, 50, 58], and domain adaptation, which adapts source knowledge to enhance target domains [10, 16, 27, 33, 61]. Besides these typical CDR works, Multi-Modal Recommendation (MMR) can also be considered a form of CDR, as different modalities can be viewed as domains due to their shared semantics [18, 28, 51, 52].

Recently, Large Language Models (LLMs) have attracted much attention for their strong performance and scalability [2, 47]. Researchers have attempted to introduce LLMs into recommender systems as well [1, 31]. However, these LLMs are Pretrained Language Models (PLMs) that are pretrained on Natural Language Processing (NLP) tasks. Therefore, researchers need to adapt these models to the recommendation domain. Under this perspective, using PLMs for recommendation can be seen as a cross-domain approach, where the source domain is the pretrained NLP domain, and the target domain is the recommendation domain.

A widely used solution is to leverage Parameter-Efficient Fine-Tuning (PEFT) techniques to perform this adaptation at affordable costs. Most works utilize Low-Rank Adaptation (LoRA) [15] for such adaptation [1, 29, 31, 56, 59]. Other techniques are also employed, such as prompt tuning [23, 25, 42, 44, 54, 57] and adapter tuning [12, 17].

2.2 Cross-Domain Sequential Recommendation

Sequential Recommendation (SR) aims to predict users' next interacted items based on their historical interaction sequences [7, 21, 45]. Similar to CDR, CDSR introduces transfer learning into SR to conduct knowledge transfer in sequential scenarios. Early CDSR

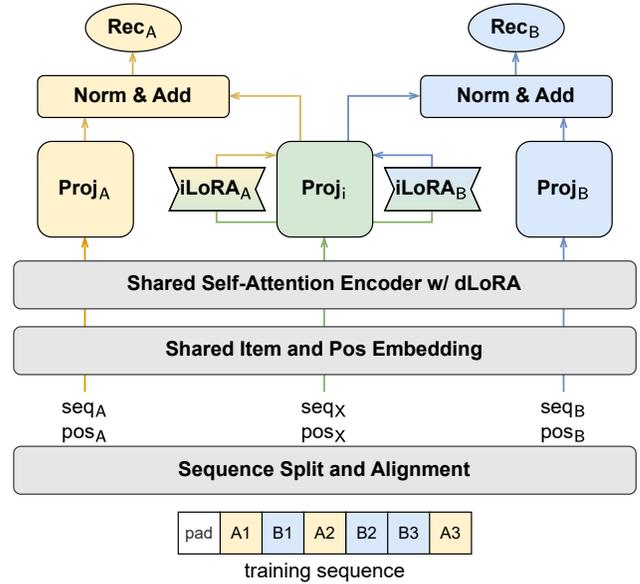


Figure 3: Processed ABXI model.

works focus on the assumption of multiple users sharing the same account [14, 38, 39, 46].

More recently, a broader concept of Cross-Domain Sequential Recommendation (CDSR) has emerged, focusing on leveraging bridging knowledge to enhance performance across domains. Depending on the type of knowledge transferred, CDSR approaches can be categorized into several types. Some studies [24, 26, 30] leverage overlapping users who have interacted in both domains as bridges to enhance performance for all users, including non-overlapping ones. Some methods [4, 5, 9, 53, 55] focus exclusively on overlapping users to strengthen their profiling in target domains. IESRec [35] leverages semantic similarities in natural language to align domains in scenarios without overlapping users. MAN [30] utilizes user groups to facilitate aligning domains. Additionally, similar to MMR, researchers have introduced multi-modal data into Sequential Recommendation (SR) to develop Multi-Modal Sequential Recommendation (MMSR) models [8, 20, 49].

3 METHODOLOGY

3.1 Problem Formulation

In this study, we focus on the dual-target CDSR task, involving two distinct domains denoted as A and B. Let a user's historical interaction sequences in domains A and B be represented as $seq_A = (i_{A1}, i_{A2}, i_{A3}, \dots, i_{An})$ and $seq_B = (i_{B1}, i_{B2}, i_{B3}, \dots, i_{Bm})$, respectively. The objective is to predict the user's next interaction items in each domain, specifically $i_{A(n+1)}$ and $i_{B(m+1)}$. The task can be formulated as:

Input: One user's domain-specific sequences, $seq_A = (i_{A1}, i_{A2}, i_{A3}, \dots, i_{An})$ and $seq_B = (i_{B1}, i_{B2}, i_{B3}, \dots, i_{Bm})$.

Output: A recommender system that estimates the probability of this users' next items, $i_{A(n+1)}$ and $i_{B(m+1)}$, to interact.

3.2 Overview

The architecture of our proposed ABXI model, as depicted in Figure 3, comprises four components: (1) sequence formulation and embedding; (2) shared self-attention encoder with domain LoRA; (3) projectors with invariant LoRA; and (4) the optimization objective.

3.3 Sequence Formulation and Embedding

Given a user’s raw training sequence, we first extract the last interaction and the first interaction to form the cross-domain sequence seq_X and the cross-domain ground truth gt_X , following the seq2seq paradigm. For simplicity, we denote the combined domain as domain X . As illustrated in Figure 2c, we then derive the domain-specific ground truths gt_A and gt_B by masking the respective domains in gt_X .

Subsequently, we create the domain-specific sequence seq_A and seq_B by masking respective domain in seq_X and re-aligning them based on gt_A and gt_B with paddings. The aligned sequences seq_A and seq_B are ensured to have the same target items, position-wise, as seq_X . In this way, the prediction mismatch issue encountered by previous works [9, 55] is addressed.

Besides, we use gt_A and gt_B solely as intermediate terms to re-align domain-specific sequences, and they are not utilized during optimization. Consequently, the ground truths of seq_X and its seq_A and seq_B are unified, which eliminates the need for standalone domain-specific recommendation loss. The position indices of each sequence are assigned separately in reverse chronological order.

Each item is then embedded into a learnable vector at length d . We initialize the item embedding table $E_I \in \mathbb{R}^{N \times d}$, where N is the total number of items, and the position embedding table $E_P \in \mathbb{R}^{L \times d}$, where L is the maximum sequence length. Besides that, we also initialize the position embedding table $E_P \in \mathbb{R}^{N \times d}$. Both embedding tables are shared across all sequences. Finally, The sequence embeddings for each sequence are obtained by adding the item and position embeddings, followed by a dropout operation to mitigate overfitting. We denote these sequence embeddings as E_X , E_A and E_B for domain X , A and B , respectively.

3.4 Low-Rank Adaptation

Many recommendation system studies [13, 19, 22, 37] have already demonstrated that index-based methods exhibit low-rank natures because of data sparsity. LoRA [15], designed as a PEFT technique, exploits similar low-rank characteristics in data to conduct efficient task adaptations [3]. We leverage LoRA to conduct efficient domain adaptation in recommendation by proposing two modules: domain LoRA (**dLoRA**), which helps shared encoders adapt to both cross-domain and domain-specific modeling; and invariant LoRA (**iLoRA**), which adapts the extracted domain-invariant knowledge into specific final recommendations. Besides, notice that incorporating dropout operators can mitigate overfitting in LoRAs [32]; Therefore, the forward pass of our proposed LoRAs yields:

$$\text{LoRA}(X) = \text{Drop} \left(\mathbf{M}_B^\uparrow \mathbf{M}_A^\downarrow X \right), \quad (1)$$

where Drop depicts dropout operator, $\mathbf{M}_A^\downarrow \in \mathbb{R}^{r \times d}$, and $\mathbf{M}_B^\uparrow \in \mathbb{R}^{d \times r}$ denote the down- and up-projection matrices with rank $r < d$. We argue that LoRA’s potential is not confined to multi-stage training

but can be effectively applied in single-stage training as well. Thus, all LoRAs in ABXI are trained together with the rest of the model.

3.5 Shared Encoder with Domain LoRAs

Inspired by the effectiveness of SASRec [21], numerous sequential recommenders adopt self-attention encoder as the backbone sequence model [6, 43, 45]. Among them, all CDSR models instantiate multiple self-attention encoders for modeling sequences from different domains [4, 5, 9, 30, 55].

We posit that a single self-attention encoder is sufficient to capture the majority of the necessary knowledge for recommendations, given the overlap in domain-invariant knowledge across domains. Consequently, we instantiate one shared encoder to all sequences. To preserve the specific uniqueness of each domain, we introduce three dLoRA modules in parallel with the encoder’s feedforward network. These dLoRAs enable the shared encoder to switch modes efficiently among the domains X , A , and B , thereby maintaining domain-specific nuances without compromising shared knowledge. The encoding process for a domain- X sequence embedding E_X is formulated as follows:

$$H_X = \text{LN} (E_X + \text{Drop} (\text{MHA} (E_X))), \quad (2)$$

$$H_X^{\text{enc}} = \text{LN} (H_X + \text{Drop} (\text{FFN} (H_X)) + \text{dLoRA}_X (H_X)), \quad (3)$$

where LN denotes LayerNorm, dLoRA_X denotes the dLoRA unit for domain X , MHA denotes the multi-head attention networks, and FFN denotes the feedforward networks. Similarly, by replacing the domain notation X with A and B in Eqs. 2 and 3, we obtain the encoded sequential representations H_A^{enc} and H_B^{enc} , respectively.

3.6 Projectors with Invariant LoRAs

To convert the encoded sequential representations into effective recommendation representations, we utilize a dedicated projector for each domain. Each domain-specific projector consists of a SwishGLU variant of MLP [47]. The structure of a projector is defined as follows:

$$\text{Proj} (X) = \text{Drop} ((\text{Swish} (XW_1) \otimes H_A^{\text{enc}} W_2) W_3), \quad (4)$$

where $W_1 \in \mathbb{R}^{d \times \frac{8}{3}d}$, $W_2 \in \mathbb{R}^{d \times \frac{8}{3}d}$ and $W_3 \in \mathbb{R}^{\frac{8}{3}d \times d}$ are learnable matrices. The Swish activation function is defined as $\text{Swish}(x) = \frac{x}{1+e^{-\beta x}}$, with β set to 1 [47].

For projectors within specific domains, we incorporate skip connections to obtain the projected domain-specific representations:

$$H_A^p = \text{LN} (H_A^{\text{enc}} + \text{Proj}_A (H_A^{\text{enc}})), \quad (5)$$

$$H_B^p = \text{LN} (H_B^{\text{enc}} + \text{Proj}_B (H_B^{\text{enc}})). \quad (6)$$

In contrast to domain-specific projectors, the invariant projector integrates two iLoRAs to adapt domain-invariant interests into final recommendations for domains A and B . The projected invariant representations are obtained as follows:

$$H_{i2A}^p = \text{LN} (H_X^{\text{enc}} + \text{Proj}_i (H_X^{\text{enc}}) + \text{iLoRA}_A (H_X^{\text{enc}})), \quad (7)$$

$$H_{i2B}^p = \text{LN} (H_X^{\text{enc}} + \text{Proj}_i (H_X^{\text{enc}}) + \text{iLoRA}_B (H_X^{\text{enc}})). \quad (8)$$

Here, Proj_i shares the same structure as Proj_A and Proj_B , while iLoRA_A and iLoRA_B are the two instantiations of iLoRA in ABXI.

The final recommendation representations are obtained by summing the projected invariant and domain-specific representations:

$$\mathbf{H}_A^{\text{rec}} = \mathbf{H}_A^{\text{P}} + \mathbf{H}_{12A}^{\text{P}}, \quad (9)$$

$$\mathbf{H}_B^{\text{rec}} = \mathbf{H}_B^{\text{P}} + \mathbf{H}_{12B}^{\text{P}}. \quad (10)$$

3.7 Optimization Objective

Previous works' adopting timestamp-guided alignment [9, 55] requires separate optimization on the cross-domain sequences. However, since our proposed task-guided alignment unites the ground truths of domain-specific and cross-domain sequences, ABXI can be optimized entirely with only one set of positive and corresponding negative samples. We split the input training sequences in the manner of the seq2seq paradigm, and randomly select N_{neg} unobserved items within the same domain for each positive sample to form the negative set. We use InfoNCE [48] to optimize ABXI. Given a recommendation representation h , we denote the embedding of the corresponding positive sample as e^+ , and the embedding set of the union of positive and negative samples as E . Therefore, the InfoNCE can thus be given as:

$$f(h) = -\log \frac{\exp(h \cdot e^+ / \tau)}{\sum_{e \in E} \exp(h \cdot e / \tau)}, \quad (11)$$

where τ denotes the temperature factor. The final loss of sequence can then be obtained as:

$$\text{Loss} = \frac{1}{|\mathbf{H}_A^{\text{rec}}|} \sum_{h \in \mathbf{H}_A^{\text{rec}}} f(h) + \frac{1}{|\mathbf{H}_B^{\text{rec}}|} \sum_{h \in \mathbf{H}_B^{\text{rec}}} f(h). \quad (12)$$

4 EXPERIMENTS

We design our experiments to answer the following research questions:

RQ1: How does ABXI perform in comparison to state-of-the-art CDSR models?

RQ2: How do the proposed task-guided alignment, projectors, iLoRA, and dLoRA of ABXI benefit its performance?

RQ3: What is the impact on ABXI's performance when replacing the proposed iLoRAs or dLoRAs with dense layers?

RQ4: How does the choice of rank hyperparameter in iLoRA and dLoRA affect the performance of ABXI?

4.1 Datasets

We conduct our experiments on three datasets derived from the Amazon review datasets¹ [40], encompassing six distinct domains: Food-Kitchen (FK), Beauty-Electronics (BE), and Movie-Book (MB). Specifically, FK includes the 'Grocery and Gourmet Food' as A and 'Home and Kitchen' as B; BE comprises 'Beauty' as A and 'Electronics' as B; MB consists of 'Movies and TV' as A and 'Books' as B.

In our preprocessing setup, each review is treated as a user interaction. We retain users who have interacted in both domains, aggregating and reordering their interactions chronologically based on the timestamps. Subsequently, we remove items that have been interacted with fewer than five times among these users. To further

¹<https://cseweb.ucsd.edu/~jmcauley/datasets/amazon/links.html>

reduce computational load, we limit each user's interaction sequence to the latest 50 interactions, following [21]. This truncation may result in some users no longer meeting the domain-overlapping criteria, necessitating a secondary filtering step to exclude these users.

The statistics of the processed datasets are summarized in Table 1. We evaluate all methods using five different random seeds to ensure the robustness and reproducibility of the results. Performance is measured using Hit Rate (HR) and Normalized Discounted Cumulative Gain (NDCG) [34] at cutoff values $K=\{5, 10\}$. For single-target models, hyperparameters are selected based on the NDCG@10 score within each domain. In contrast, hyperparameters are optimized for dual-target models based on the aggregate NDCG@10 scores across both domains.

4.2 Baselines

We compare ABXI with several baseline models, categorized into four types:

ST-SDSR Single-Target Single-Domain Sequential Recommenders: SASRec-1 and BERT4Rec-1.

DT-SDSR Dual-Target Single-Domain Sequential Recommender: SASRec-2 and BERT4Rec-2.

ST-CDSR Single-Target Cross-Domain Sequential Recommender: CD-SASRec, CD-ASR and MGCL.

DT-CDSR Dual-Target Cross-Domain Sequential Recommender: C²DSR and DREAM.

These baseline models are described as follows:

- **SASRec** [21] utilizes the self-attention encoder to generate sequential representations. We implement two versions: ST-SASRec-1 and SASRec-2, corresponding to ST and DT settings, respectively. Specifically, losses in SASRec-2 are calculated separately within each domain and then summed.
- **BERT4Rec** [45] introduce Cloze objectives on top of SASRec. Similarly, we use two versions: BERT4Rec-1 and BERT4Rec-2.
- **CD-SASRec** [5] aggregates the encoded source-domain sequences into the target-domain encoding using two self-attention encoders.
- **CD-ASR** [4] fuses source and target domain sequences encoded by separate self-attention encoders.
- **C²DSR** [9] instantiate difference set of graphical and self-attention encoder to encode cross-domain and domain-specific sequences, leveraging augmentation for contrastive learning.
- **MGCL** [53] integrates graphical and sequential information under different views and strengthens the profiling via user-to-user contrastive learning on views.
- **DREAM** [55] employs separate self-attention encoders for cross-domain and domain-specific sequences, incorporating specific-to-cross knowledge transferring and a similar user-to-user contrastive learning on domains.

4.3 Implementation Details

We adopt the leave-one-out strategy commonly used in SR. Specifically, we remove the last two interactions from each user sequence to serve as the ground truths for validation and testing. Evaluation metrics are computed separately for each domain. Besides basic statistics, Table 1 summarizes the number of validation and testing

Table 1: Statistics of CDSR Datasets.

Datasets	FK		BE		MB	
	Food (A)	Kitchen (B)	Beauty (A)	Electronics (B)	Movie (A)	Book (B)
Users	7,144		4,474		28,350	
Items	11,837	16,258	10,379	14,188	35,712	90,958
Interactions	83,663	89,885	50,329	63,800	347,654	403,147
Val. GT	2,837	4,307	2,086	2,388	11,728	16,622
Test GT	2,419	4,725	1,875	2,599	10,935	17,415
A → B transitions	30,308	-	17,841	-	108,318	-
B → A transitions	-	29,407	-	17,888	-	105,696

Table 2: Hyperparameters selection.

Hyperparameter	Value
Embedding dimension d	256
# self-attention layer	1
dropout rate	0.3
# negative sample N_{neg}	128
Optimizer	AdamW
Temperature τ	0.75
Max epoch	500
Warm-up epoch	5
Learning rate	{1e-3, 1e-4}
Weight decay	{5, 2, 1} × {1e1, 1e0, 1e-1, 1e-2, 1e-3}, 0
Learning rate decay	×0.3162, after 30 stable epochs
Random seed	{3407, 0, 1, 2, 3}

ground truths per domain, as well as the counts of cross-domain item-to-item transitions within all sequences (i.e., $A \rightarrow B$ and $B \rightarrow A$). Hyperparameters not specified elsewhere are listed in Table 2.

Our experimental setup treats CDSR as SDRS with added side domain information. By comparing both types of models under identical conditions, we aim to assess the performance improvements fairly through CDSR. If new CDSR models do not outperform classic SDRS models within this setting, their practical utility may be limited.

4.4 Overall Performance Comparison (RQ1)

To evaluate ABXI and address RQ1, we compare our model with state-of-the-art CDSR models and other baseline methods. For each dataset, we designate the domain with better performance metrics as the ‘easy’ domain and the other as the ‘hard’ domain. The overall results are reported in Table 3.

Across all datasets, ABXI significantly outperforms all baseline models on all evaluation metrics, with statistical significance ($p < 0.01$). This improvement is evident in both domains, indicating that a single sequence model augmented with auxiliary modules is sufficient to capture both cross-domain and domain-specific knowledge.

In contrast, other DT-CDSR models do not demonstrate substantial improvements over their SDRS counterparts. Only in the Movie domain do we observe that DREAM outperforms the SDRS models.

This can be attributed to the prediction mismatches in timestamp-guided alignment within these DT-CDSR models, causing them to make predictions in incorrect domains. Conversely, ST- and DT-SDSR models do not encounter such issues, as they do not require domain-specific sequences.

Furthermore, the ST-CDSR models demonstrate notably poor performance compared to all other types of models. This underperformance stems from their training from scratch without supervision in the source domains, which hinders the extraction of domain-invariant information. Without supervision from the source domain, these models cannot effectively distinguish domain-invariant knowledge from domain-specific knowledge. This results in source-specific information, which should have been filtered out, being introduced into the target domain.

4.5 Ablation Studies (RQ2)

To further investigate the contribution of each proposed component, we designed five ablation variants: V_1 removes all dLoRAs; V_2 removes all projectors; V_3 removes all iLoRAs; V_4 removes all dLoRAs, iLoRAs and projectors; V_5 replaces the task-guided alignment with the timestamp-guided alignment. The results of these ablation studies are reported in the middle section of Table 4.

Compared to all ablation variants, ABXI achieves the best performance. Notably, we have the following observations.

Projectors contribute the most among all trainable modules. Due to their substantial number of learnable parameters, projectors provide sufficient capacity to transform encoded sequential representations into recommendation representations.

The performance gaps between V_1 , V_3 and ABXI, as well as between V_2 and V_4 , indicate that iLoRA and dLoRA effectively enhance performance with few additional parameters.

The performance improvements of ABXI over V_1 and V_3 , as well as of V_2 over V_4 , demonstrate that iLoRA and dLoRA effectively enhance performance with a small number of additional parameters.

Among all variants, reverting to timestamp-guided alignment in V_5 results in the most significant performance degradation. Since ABXI does not employ specialized domain-specific ground truths under timestamp-guided alignment, it suffers severely from the prediction mismatch issue.

Table 3: Recommendation performance (RQ1). The best and the runner-up are highlighted in bold and underlined respectively.

Type	Methods	Food				Kitchen			
		HR@5	HR@10	NDCG@5	NDCG@10	HR@5	HR@10	NDCG@5	NDCG@10
ST	SASRec-1	0.1867±0.0041	0.2530±0.0037	0.1309±0.0033	0.1524±0.0032	0.1213±0.0029	0.1786±0.0019	0.0822±0.0014	0.1006±0.0014
	BERT4Rec-1	0.1812±0.0030	0.2482±0.0038	0.1244±0.0021	0.1461±0.0022	0.1099±0.0042	0.1652±0.0059	0.0736±0.0032	0.0914±0.0036
	CD-SASRec	0.1641±0.0163	0.2277±0.0135	0.1119±0.0131	0.1326±0.0120	0.0993±0.0060	0.1594±0.0082	0.0645±0.0044	0.0837±0.0045
	CD-ASR	0.1892±0.0022	0.2590±0.0050	0.1287±0.0025	0.1513±0.0033	0.1146±0.0016	0.1727±0.0023	0.0760±0.0017	0.0946±0.0018
	MGCL	0.1753±0.0061	0.2465±0.0052	0.1191±0.0026	0.1423±0.0025	0.1173±0.0073	0.1803±0.0111	0.0782±0.0046	0.0985±0.0059
DT	SASRec-2	<u>0.2269</u> ±0.0021	<u>0.2857</u> ±0.0017	<u>0.1568</u> ±0.0018	<u>0.1760</u> ±0.0024	0.1458±0.0030	<u>0.2105</u> ±0.0031	<u>0.0988</u> ±0.0018	<u>0.1196</u> ±0.0018
	BERT4Rec-2	0.2157±0.0046	0.2827±0.0042	0.1475±0.0054	0.1693±0.0041	0.1366±0.0051	0.2047±0.0050	0.0904±0.0039	0.1123±0.0038
	C ² DSR	0.2071±0.0074	0.2669±0.0053	0.1453±0.0028	0.1647±0.0029	0.1244±0.0037	0.1841±0.0069	0.0850±0.0024	0.1043±0.0033
	DREAM	0.1947±0.0041	0.2521±0.0090	0.1369±0.0022	0.1555±0.0038	0.1244±0.0039	0.1790±0.0060	0.0830±0.0031	0.1007±0.0037
	ABXI	0.2499 ±0.0037	0.3187 ±0.0038	0.1752 ±0.0035	0.1977 ±0.0033	0.1757 ±0.0039	0.2439 ±0.0029	0.1205 ±0.0017	0.1425 ±0.0009
		Beauty				Electronics			
		HR@5	HR@10	NDCG@5	NDCG@10	HR@5	HR@10	NDCG@5	NDCG@10
ST	SASRec-1	0.1616±0.0021	0.2286±0.0041	0.1130±0.0020	0.1346±0.0027	0.1172±0.0026	0.1711±0.0031	0.0803±0.0009	0.0976±0.0010
	BERT4Rec-1	0.1648±0.0031	0.2354±0.0037	0.1141±0.0034	0.1368±0.0037	0.1307±0.0020	0.1870±0.0016	0.0901±0.0016	0.1082±0.0013
	CD-SASRec	0.1303±0.0104	0.2016±0.0128	0.0887±0.0081	0.1116±0.0083	0.1131±0.0072	0.1645±0.0104	0.0766±0.0050	0.0932±0.0057
	CD-ASR	0.1587±0.0042	0.2389±0.0042	0.1082±0.0035	0.1341±0.0033	0.1258±0.0029	0.1842±0.0020	0.0856±0.0027	0.1044±0.0023
	MGCL	0.1400±0.0076	0.2096±0.0090	0.0952±0.0056	0.1176±0.0060	0.1272±0.0033	0.1828±0.0055	0.0863±0.0027	0.1043±0.0031
DT	SASRec-2	<u>0.2371</u> ±0.0062	<u>0.3210</u> ±0.0037	<u>0.1678</u> ±0.0028	<u>0.1949</u> ±0.0020	0.1353±0.0032	0.2041±0.0038	0.0901±0.0027	0.1124±0.0021
	BERT4Rec-2	0.1810±0.0060	0.2794±0.0063	0.1228±0.0046	0.1545±0.0048	<u>0.1467</u> ±0.0029	<u>0.2148</u> ±0.0026	<u>0.1002</u> ±0.0025	<u>0.1221</u> ±0.0028
	C ² DSR	0.1927±0.0090	0.2785±0.0122	0.1310±0.0051	0.1588±0.0058	0.1361±0.0052	0.1905±0.0060	0.0930±0.0033	0.1105±0.0030
	DREAM	0.2007±0.0077	0.2918±0.0067	0.1362±0.0045	0.1656±0.0035	0.1164±0.0065	0.1733±0.0071	0.0794±0.0043	0.0977±0.0043
	ABXI	0.2722 ±0.0024	0.3722 ±0.0025	0.1862 ±0.0030	0.2186 ±0.0025	0.1642 ±0.0044	0.2389 ±0.0061	0.1138 ±0.0036	0.1377 ±0.0041
		Movie				Book			
		HR@5	HR@10	NDCG@5	NDCG@10	HR@5	HR@10	NDCG@5	NDCG@10
ST	SASRec-1	0.2258±0.0031	0.2961±0.0037	0.1647±0.0025	0.1874±0.0027	0.1357±0.0029	0.1789±0.0033	0.1007±0.0022	0.1147±0.0023
	BERT4Rec-1	0.2155±0.0021	0.2907±0.0019	0.1546±0.0016	0.1789±0.0015	<u>0.1429</u> ±0.0020	<u>0.1874</u> ±0.0025	<u>0.1069</u> ±0.0009	<u>0.1212</u> ±0.0011
	CD-SASRec	0.2187±0.0065	0.2944±0.0074	0.1555±0.0044	0.1800±0.0047	0.1413±0.0081	0.1829±0.0082	0.1045±0.0065	0.1179±0.0064
	CD-ASR	0.2088±0.0011	0.2819±0.0015	0.1492±0.0015	0.1728±0.0017	0.1294±0.0015	0.1742±0.0014	0.0925±0.0014	0.1070±0.0013
	MGCL	0.1980±0.0084	0.2662±0.0078	0.1416±0.0083	0.1636±0.0080	0.1163±0.0011	0.1533±0.0016	0.0866±0.0009	0.0985±0.0010
DT	SASRec-2	0.2150±0.0026	0.2837±0.0013	0.1557±0.0013	0.1779±0.0011	0.1251±0.0014	0.1671±0.0016	0.0931±0.0011	0.1066±0.0013
	BERT4Rec-2	0.2064±0.0022	0.2790±0.0028	0.1476±0.0014	0.1710±0.0017	0.1399±0.0011	0.1854±0.0016	0.1038±0.0004	0.1185±0.0005
	C ² DSR	0.2135±0.0049	0.2813±0.0042	0.1551±0.0036	0.1770±0.0034	0.1155±0.0053	0.1538±0.0063	0.0863±0.0037	0.0987±0.0039
	DREAM	<u>0.2308</u> ±0.0073	<u>0.3060</u> ±0.0084	<u>0.1665</u> ±0.0058	<u>0.1908</u> ±0.0062	0.1299±0.0058	0.1766±0.0076	0.0938±0.0043	0.1089±0.0047
	ABXI	0.2790 ±0.0030	0.3590 ±0.0015	0.2078 ±0.0016	0.2335 ±0.0014	0.1901 ±0.0015	0.2472 ±0.0015	0.1429 ±0.0009	0.1613 ±0.0007

Table 4: Ablation studies in NDCG@10. (RQ2, 3)

	iLoRA	Proj	dLoRA	Alignment	Food	Kitchen	Beauty	Electronics
ABXI	✓	✓	✓	task	0.1977 ±0.0033	0.1425 ±0.0009	<u>0.2186</u> ±0.0025	0.1377 ±0.0041
V ₁	✓	✓	-	task	0.1977 ±0.0026	<u>0.1417</u> ±0.0017	0.2151±0.0020	0.1353±0.0014
V ₂	✓	-	✓	task	0.1938±0.0013	0.1381±0.0015	0.2092±0.0015	0.1345±0.0039
V ₃	-	✓	✓	task	0.1970±0.0027	0.1402±0.0011	0.2173±0.0028	0.1349±0.0026
V ₄	-	-	-	task	0.1923±0.0034	0.1327±0.0018	0.2062±0.0024	<u>0.1366</u> ±0.0026
V ₅	✓	✓	✓	timestamp	0.1753±0.0028	0.1305±0.0032	0.2011±0.0082	0.1198±0.0017
ABXI-i3	3×Proj	✓	✓	task	0.1954±0.0025	0.1415±0.0022	0.2185±0.0022	0.1348±0.0032
ABXI-i2	2×Proj	✓	✓	task	0.1953±0.0037	0.1413±0.0020	0.2191 ±0.0042	<u>0.1366</u> ±0.0034
ABXI-d	✓	✓	3×encoders	task	0.1976±0.0016	0.1405±0.0018	0.2114±0.0042	0.1329±0.0031

4.6 Effectiveness of LoRA (RQ3)

LoRA is renowned for its efficient adaptation. However, in the context of non-LLM CDSR, we need to determine whether this efficiency translates into effectiveness. Therefore, we design the following three variants: **ABXI-i3**: replace $i\text{LoRA}_A$ and $i\text{LoRA}_B$ with X-to-A and X-to-A projectors; **ABXI-i2**: remove the Proj_i in ABXI-i3; **ABXI-d**: replace the shared encoder with dLoRAs with three self-attention encoders for domains X , A and B , respectively. The results are reported in the bottom section of Table 4.

Among them, ABXI-i2 achieves performance closely matching that of ABXI. This suggests redundancy between the X-to-A and X-to-A projectors, which indicates that they can be functionally replaced by the combination of the Proj_i and two $i\text{LoRA}$ modules.

ABXI-i3 utilizes the shared Proj_i along with specialized X-to-A and X-to-A projectors. However, this configuration does not yield significant performance improvements. The introduction of these projectors substantially increases the capacity for cross-to-specific transformations, which diminishes ABXI's reliance on the Proj_i and impedes its ability to extract domain-invariant knowledge.

ABXI-d degrades performance by instantiating three encoders for each domain instead of the shared one with dLoRAs. We attribute this degradation to our task-guided alignment. Since we unify the downstream recommendation tasks, we do not need to employ additional domain-specific ground truths like $C^2\text{DSR}$ [9] and DREAM [55]. Therefore, these domain-specific encoders cannot achieve the same performance as they do in those baselines.

4.7 Rank Analysis in LoRA (RQ4)

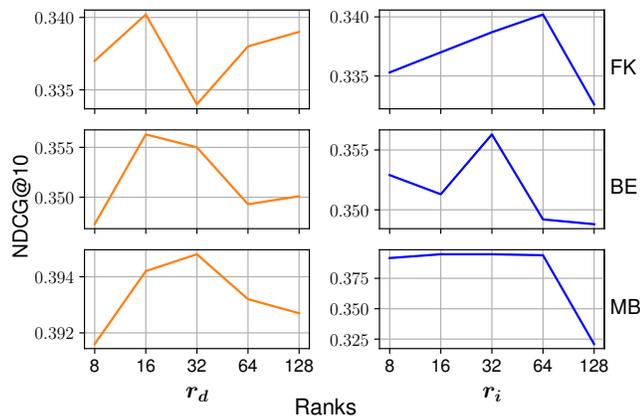


Figure 4: Performance of all domains under different values of ranks of dLoRA (r_d) and iLoRA (r_i) in NDCG@10 (RQ4).

This subsection investigates the impact of different rank values r_d for dLoRA and r_i for iLoRA on the ABXI's performance. We vary these ranks over $\{8, 16, 32, 64, 128\}$ and present the results in Figure 4.

From Figure 4, we observe a general tendency across all six domains: both large and small ranks tend to yield poorer results, although some fluctuations are evident. This behavior is due to the limited capacity of LoRA modules within the one-stage training

framework of ABXI. To achieve optimal performance, we select ($r_d = 16, r_i = 64$) for FK, ($r_d = 16, r_i = 32$) for BE, and ($r_d = 32, r_i = 32$) for MB.

5 CONCLUSION

In this paper, we propose ABXI, a novel CDSR model. Specifically, ABXI addresses the prediction mismatch issue by integrating task-guided alignment to unify cross-domain and domain-specific optimizations. ABXI leverages domain projectors equipped with invariant LoRA modules to extract and adapt domain-invariant interests, enabling efficient and effective knowledge transfer. Furthermore, ABXI employs a single shared encoder with domain LoRA to conduct efficient encoding. Extensive experimental results on three datasets demonstrated that ABXI significantly outperforms state-of-the-art CDSR models by a large margin. Ablation studies confirmed the effectiveness of each component, highlighting the importance of task-guided alignment and invariant interest extraction. For future work, we aim to further explore domain-invariant dynamics to achieve more accurate disentanglement.

REFERENCES

- [1] Arkadeep Acharya, Brijraj Singh, and Naoyuki Onoe. 2023. Llm based generation of item-description for recommendation system. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1204–1207.
- [2] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [3] Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. 2020. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255* (2020).
- [4] Nawaf Alharbi and Doina Caragea. 2021. Cross-domain attentive sequential recommendations based on general and current user preferences (cd-asr). In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*. 48–55.
- [5] Nawaf Alharbi and Doina Caragea. 2022. Cross-domain self-attentive sequential recommendations. In *Proceedings of International Conference on Data Science and Applications: ICDSA 2021, Volume 2*. Springer, 601–614.
- [6] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1007–1014.
- [7] Qingtian Bian, Jiaxing Xu, Hui Fang, and Yiping Ke. 2023. CPMR: Context-Aware Incremental Sequential Recommendation with Pseudo-Multi-Task Learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 120–130.
- [8] Shuqing Bian, Xingyu Pan, Wayne Xin Zhao, Jinpeng Wang, Chuyuan Wang, and Ji-Rong Wen. 2023. Multi-modal Mixture of Experts Representation Learning for Sequential Recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 110–119.
- [9] Jiangxia Cao, Xin Cong, Jiawei Sheng, Tingwen Liu, and Bin Wang. 2022. Contrastive Cross-Domain Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 138–147.
- [10] Jiangxia Cao, Xixun Lin, Xin Cong, Jing Ya, Tingwen Liu, and Bin Wang. 2022. Disencdr: Learning disentangled representations for cross-domain recommendation. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*. 267–277.
- [11] Shu Chen, Zitao Xu, Weike Pan, Qiang Yang, and Zhong Ming. 2024. A Survey on Cross-Domain Sequential Recommendation. *arXiv preprint arXiv:2401.04971* (2024).
- [12] Junchen Fu, Fajie Yuan, Yu Song, Zheng Yuan, Mingyue Cheng, Shenghui Cheng, Jiaqi Zhang, Jie Wang, and Yunzhu Pan. 2024. Exploring adapter-based transfer learning for recommender systems: Empirical studies and practical insights. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 208–217.
- [13] Guibing Guo, Jie Zhang, and Neil Yorke-Smith. 2015. Trustsvd: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 29.

- [14] Lei Guo, Li Tang, Tong Chen, Lei Zhu, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2021. DA-GCN: A domain-aware attentive graph convolution network for shared-account cross-domain sequential recommendation. *arXiv preprint arXiv:2105.03300* (2021).
- [15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).
- [16] Guangneng Hu, Yu Zhang, and Qiang Yang. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 667–676.
- [17] Jun Hu, Wenwen Xia, Xiaolu Zhang, Chilin Fu, Weichang Wu, Zhaoxin Huan, Ang Li, Zuoli Tang, and Jun Zhou. 2024. Enhancing sequential recommendation via llm-based semantic embedding learning. In *Companion Proceedings of the ACM on Web Conference 2024*. 103–111.
- [18] Zhenhua Huang, Xin Xu, Juan Ni, Honghao Zhu, and Cheng Wang. 2019. Multi-modal representation learning for recommendation in Internet of Things. *IEEE Internet of Things Journal* 6, 6 (2019), 10675–10685.
- [19] Mohsen Jamali and Martin Ester. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*. 135–142.
- [20] Wei Ji, Xiangyan Liu, An Zhang, Yinwei Wei, Yongxin Ni, and Xiang Wang. 2023. Online distillation-enhanced multi-modal transformer for sequential recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia*. 955–965.
- [21] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [22] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 426–434.
- [23] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691* (2021).
- [24] Chenglin Li, Mingjun Zhao, Huanming Zhang, Chenyun Yu, Lei Cheng, Guoqiang Shu, Beibei Kong, and Di Niu. 2022. RecGURU: Adversarial learning of generalized user representations for cross-domain recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*. 571–581.
- [25] Lei Li, Yongfeng Zhang, and Li Chen. 2023. Prompt distillation for efficient llm-based recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 1348–1357.
- [26] Pan Li, Zhichao Jiang, Maofei Que, Yao Hu, and Alexander Tuzhilin. 2021. Dual attentive sequential learning for cross-domain click-through rate prediction. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3172–3180.
- [27] Pan Li and Alexander Tuzhilin. 2020. Dtdcdr: Deep dual transfer cross domain recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 331–339.
- [28] Xiaopeng Li, Fan Yan, Xiangyu Zhao, Yichao Wang, Bo Chen, Hui Feng Guo, and Ruiming Tang. 2023. Hamur: Hyper adapter for multi-domain recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 1268–1277.
- [29] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. Llara: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1785–1795.
- [30] Guanyu Lin, Chen Gao, Yu Zheng, Jianxin Chang, Yanan Niu, Yang Song, Kun Gai, Zhiheng Li, Depeng Jin, Yong Li, et al. 2024. Mixed attention network for cross-domain sequential recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 405–413.
- [31] Xinyu Lin, Wenjie Wang, Yongqi Li, Shuo Yang, Fuli Feng, Yinwei Wei, and Tat-Seng Chua. 2024. Data-efficient Fine-tuning for LLM-based Recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 365–374.
- [32] Yang Lin, Xinyu Ma, Xu Chu, Yujie Jin, Zhibang Yang, Yasha Wang, and Hong Mei. 2024. Lora dropout as a sparsity regularizer for overfitting control. *arXiv preprint arXiv:2404.09610* (2024).
- [33] Meng Liu, Jianjun Li, Guohui Li, and Peng Pan. 2020. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 885–894.
- [34] Tie-Yan Liu et al. 2009. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* 3, 3 (2009), 225–331.
- [35] Weiming Liu, Xiaolin Zheng, Chaochao Chen, Jiajie Su, Xinting Liao, Mengling Hu, and Yanhao Tan. 2023. Joint internal multi-interest exploration and external domain alignment for cross domain sequential recommendation. In *Proceedings of the ACM Web Conference 2023*. 383–394.
- [36] Haokai Ma, Ruobing Xie, Lei Meng, Xin Chen, Xu Zhang, Leyu Lin, and Jie Zhou. 2024. Triple sequence learning for cross-domain recommendation. *ACM Transactions on Information Systems* 42, 4 (2024), 1–29.
- [37] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and knowledge management*. 931–940.
- [38] Muyang Ma, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Lifan Zhao, Peiyu Liu, Jun Ma, and Maarten de Rijke. 2022. Mixed information flow for cross-domain sequential recommendations. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 4 (2022), 1–32.
- [39] Muyang Ma, Pengjie Ren, Yujie Lin, Zhumin Chen, Jun Ma, and Maarten de Rijke. 2019. π -net: A parallel information-sharing network for shared-account cross-domain sequential recommendations. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 685–694.
- [40] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 43–52.
- [41] Chung Park, Taesan Kim, Taekyoon Choi, Junui Hong, Yelim Yu, Mincheol Cho, Kyunam Lee, Sungil Ryu, Hyungjun Yoon, Minsung Choi, et al. 2023. Cracking the Code of Negative Transfer: A Cooperative Game Theoretic Approach for Cross-Domain Sequential Recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 2024–2033.
- [42] Gustavo Penha and Claudia Hauff. 2020. What does bert know about books, movies and music? probing bert for conversational recommendation. In *Proceedings of the 14th ACM conference on recommender systems*. 388–397.
- [43] Aleksandr Petrov and Craig Macdonald. 2022. gSASRec: Reducing Overconfidence in Sequential Recommendation Trained with Negative Sampling. In *Proceedings of the 17th ACM Conference on Recommender Systems*.
- [44] Tianshu Shen, Jiuru Li, Mohamed Reda Bouadjenek, Zheda Mai, and Scott Sanner. 2023. Towards understanding and mitigating unintended biases in language model-driven conversational recommendation. *Information Processing & Management* 60, 1 (2023), 103139.
- [45] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM International conference on Information and Knowledge Management*. 1441–1450.
- [46] Wenchao Sun, Muyang Ma, Pengjie Ren, Yujie Lin, Zhumin Chen, Zhaochun Ren, Jun Ma, and Maarten De Rijke. 2021. Parallel split-join networks for shared account cross-domain sequential recommendations. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [47] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmina Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [48] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR* abs/1807.03748 (2018). arXiv:1807.03748 <http://arxiv.org/abs/1807.03748>
- [49] Jinpeng Wang, Ziyun Zeng, Yunxiao Wang, Yuting Wang, Xingyu Lu, Tianxiang Li, Jun Yuan, Rui Zhang, Hai-Tao Zheng, and Shu-Tao Xia. 2023. MISSRec: Pre-training and transferring multi-modal interest-aware sequence representation for recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia*. 6548–6557.
- [50] Tianxin Wang, Fuzhen Zhuang, Zhiqiang Zhang, Daixin Wang, Jun Zhou, and Qing He. 2021. Low-dimensional alignment for cross-domain recommendation. In *Proceedings of the 30th ACM international conference on information & knowledge management*. 3508–3512.
- [51] Wei Wei, Chao Huang, Lianghao Xia, and Chuxu Zhang. 2023. Multi-modal self-supervised learning for recommendation. In *Proceedings of the ACM Web Conference 2023*. 790–800.
- [52] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM international conference on multimedia*. 1437–1445.
- [53] Zitao Xu, WeiKe Pan, and Zhong Ming. 2023. A Multi-view Graph Contrastive Learning Framework for Cross-Domain Sequential Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 491–501.
- [54] Bowen Yang, Cong Han, Yu Li, Lei Zuo, and Zhou Yu. 2021. Improving Conversational Recommendation Systems' Quality with Context-Aware Item Meta Information. *arXiv preprint arXiv:2112.08140* (2021).
- [55] Xiaoxin Ye, Yun Li, and Lina Yao. 2023. DREAM: Decoupled Representation via Extraction Attention Module and Supervised Contrastive Learning for Cross-Domain Sequential Recommender. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 479–490.
- [56] Bin Yin, Junjie Xie, Yu Qin, Zixiang Ding, Zhichao Feng, Xiang Li, and Wei Lin. 2023. Heterogeneous knowledge fusion: A novel approach for personalized recommendation via llm. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 599–601.

1045	[57]	Zizhuo Zhang and Bang Wang. 2023. Prompt learning for news recommendation. In <i>Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval</i> . 227–237.	[60]	Xiaolin Zheng, Jiajie Su, Weiming Liu, and Chaochao Chen. 2022. DDGHM: dual dynamic graph with hybrid metric training for cross-domain sequential recommendation. In <i>Proceedings of the 30th ACM International Conference on Multimedia</i> . 471–481.	1103
1046					1104
1047	[58]	Chuang Zhao, Hongke Zhao, Ming He, Jian Zhang, and Jianping Fan. 2023. Cross-domain recommendation via user interest alignment. In <i>Proceedings of the ACM Web Conference 2023</i> . 887–896.	[61]	Yongchun Zhu, Zhenwei Tang, Yudan Liu, Fuzhen Zhuang, Ruobing Xie, Xu Zhang, Leyu Lin, and Qing He. 2022. Personalized transfer of user preferences for cross-domain recommendation. In <i>Proceedings of the fifteenth ACM international conference on web search and data mining</i> . 1507–1515.	1105
1048					1106
1049	[59]	Jujia Zhao, Wenjie Wang, Chen Xu, Zhaochun Ren, See-Kiong Ng, and Tat-Seng Chua. 2024. Llm-based federated recommendation. <i>arXiv preprint arXiv:2402.09959</i> (2024).			1107
1050					1108
1051					1109
1052					1110
1053					1111
1054					1112
1055					1113
1056					1114
1057					1115
1058					1116
1059					1117
1060					1118
1061					1119
1062					1120
1063					1121
1064					1122
1065					1123
1066					1124
1067					1125
1068					1126
1069					1127
1070					1128
1071					1129
1072					1130
1073					1131
1074					1132
1075					1133
1076					1134
1077					1135
1078					1136
1079					1137
1080					1138
1081					1139
1082					1140
1083					1141
1084					1142
1085					1143
1086					1144
1087					1145
1088					1146
1089					1147
1090					1148
1091					1149
1092					1150
1093					1151
1094					1152
1095					1153
1096					1154
1097					1155
1098					1156
1099					1157
1100					1158
1101					1159
1102					1160