

# CD-CDR: Conditional Diffusion-based Item Generation for Cross-Domain Recommendation

Anonymous Author(s)

## ABSTRACT

Cross-domain recommendation (CDR) has emerged as a promising direction for expanding the applicability of recommendation systems. Recent advances in CDR have demonstrated the effectiveness of the unified distribution paradigm, which leverages shared distributions to transfer knowledge across domains and employ domain-specific adapters for targeted recommendations. While this well-designed paradigm shows promising performance, existing methods require extra supervision signals (e.g. contrastive learning on domain-masked embeddings) to maintain unified distributions across domains, leading to an inherent trade-off between unified objectives and domain-specific preference modeling.

To address these limitations, we propose **CD-CDR** (Conditional Diffusion-CDR), a novel approach that leverages a shared conditional diffusion model to learn unified item distributions and facilitate knowledge transfer across domains. The key insight is to utilize the powerful generative capabilities of diffusion models to learn a shared distribution, while naturally incorporating domain-specific characteristics through conditional generation. This design enables CD-CDR to replace traditional adapters with generation conditions as an integral part of the distribution model, thereby eliminating extra supervision signals and fundamentally resolving the trade-off between unified and domain-specific objectives. Extensive experiments on six domain pairs from two real-world datasets demonstrate that CD-CDR significantly outperforms existing methods for both normal and cold-start settings. To the best of our knowledge, this is the first work to explore the unified distribution paradigm in CDR using conditional diffusion models.

## CCS CONCEPTS

• Information systems → Recommender systems.

## KEYWORDS

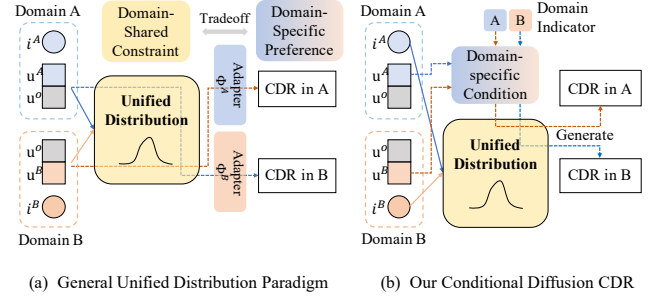
Recommender Systems, Cross-Domain Recommendation, User Modeling, Diffusion Models

## ACM Reference Format:

Anonymous Author(s). 2025. CD-CDR: Conditional Diffusion-based Item Generation for Cross-Domain Recommendation. In *SIGIR '25, July 13–July 18 2025, Italy*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

In recent years, recommendation systems have played a pivotal role in improving user experiences across various platforms by providing personalized content. As the field evolves, Cross-Domain Recommendation (CDR) [31, 36] has emerged as a promising research direction, aiming to leverage information from different domains to improve recommendation performance. There are two types of CDR tasks based on whether the recommended items are in the same domain as the user: intra-domain and inter-domain



**Figure 1: Illustration of the unified distribution paradigm for cross-domain recommendation models and our proposed CD-CDR, where  $u^o$  denotes overlap users across two domains.**

recommendation [1]. Intra-domain CDR focuses on recommending items to users in the same domain, where the primary challenge lies in effectively transferring knowledge from other domains to enhance recommendation quality in the target domain. In contrast, inter-domain recommendation aims to recommend items from one domain to users from other domains, which essentially becomes a cold-start problem since these users have no interaction history in the target item domain. Both tasks are fundamental to advancing recommendation systems, particularly in scenarios where user data is sparse or distributed across different domains.

To tackle these challenges, existing work has developed two mainstream paradigms: the **embedding & mapping paradigm** [11, 15, 21, 27, 38] and the **unified distribution paradigm** [2, 4, 24, 29]. The embedding & mapping paradigm learns user mapping functions from overlapping users to transfer information across domains. While extensively explored through various techniques like meta-learning [38], variational autoencoders [32], and recently diffusion models [27], this paradigm faces inherent limitations. As noted by Cao et al. [4], independently pre-trained representations can lead to domain-specific biases, and the assumption of shared mapping functions across all users may not hold. In contrast, the unified distribution paradigm, as shown in the left part of Figure 1, aims to learn domain-shared user/item representations by considering interactions from both domains simultaneously, and recommendations in each domain are provided by domain-specific adapters. This paradigm provides an end-to-end framework for CDR tasks without assumptions about users in different domains and can be easily extended to multiple domains without strict constraints about the source and the target domains. Early-stage research such as CMF [24] directly mixes different domain data with the same model. Recent work has employed various techniques to simultaneously model the shared knowledge between domains and the behavioral shift across domains, such as contrastive learning [2, 29], adversarial learning [6] and information bottleneck [4]. These approaches have shown promising results in capturing shared knowledge between domains. However, the unified distribution paradigm is usually

constrained by extra supervision signals to capture the distribution, causing a trade-off between unified and domain-specific objectives.

In this paper, we aim to **address inter- and intra-domain CDR tasks simultaneously with the unified distribution paradigm based on the conditional diffusion model**. Recent advances in generative models, particularly diffusion models [7] provide a possible new way to model the unified distributions across different domains in inter- and intra-domain CDR tasks. Diffusion models have demonstrated remarkable capabilities in modeling complex distributions and generating high-quality samples [17, 23]. Specifically, conditional diffusion models [8, 17, 20] offer a natural way to describe unified distributions across domains while maintaining domain-specific characteristics by conditioning for distribution generations. Furthermore, these models eliminate the need for extra objectives to avoid modeling domain-specific behaviors in the unified distribution, as the conditional input inherently serves two functions: providing historical context and acting as domain-specific adapters. However, applying diffusion models to the unified distribution modeling of CDR presents several unique challenges: First, unlike traditional applications where diffusion models generate data in continuous space, recommendation scenarios involve discrete user-item interactions, requiring careful adaptation of the diffusion process. Second, while our goal is to learn a unified distribution across domains, we must simultaneously preserve domain-specific features to ensure personalized and relevant recommendations. Third, the model needs to effectively adapt the unified distributions to different domains, particularly for inter-domain tasks where target domain information is limited.

To leverage the advantages of diffusion models in CDR while addressing the aforementioned challenges, we propose CD-CDR, an innovative method that adopts a conditional diffusion model to create a unified item distribution across different domains with cross-domain conditional generation. Our proposed model integrates historical interactions from both the source and target domains, balancing their influences to generate potential items of interest based on the combined behavior patterns. At the core of CD-CDR is a shared diffusion model that captures item distributions across domains, facilitating the seamless transfer of user preferences. This method leverages two key strengths of diffusion models: their ability to handle noise which reduces the impact of irrelevant domain information, and their capacity to establish consistent item distributions across domains. As a result, CD-CDR effectively recommends items for cold-start users in the target domain by incorporating historical data from the source domain. To validate the efficacy of our proposed model, we conducted extensive experiments on six domain pairs from two real-world datasets. The experimental results demonstrate the superiority of CD-CDR over existing cross-domain and single-domain methods on both inter- and intra-domain tasks, highlighting its ability to improve recommendation accuracy in CDR scenarios significantly.

Our contributions can be summarized as follows:

- To the best of our knowledge, this is the first attempt to explore conditional diffusion models with the unified distribution paradigm for CDR scenarios.
- We introduce a novel CDR approach, CD-CDR, that naturally incorporates domain-specific characteristics through

conditional item generation from a shared diffusion model depicting a domain-shared item distribution.

- Comprehensive experimental evaluations validate that CD-CDR outperforms SOTA methods from both existing CDR paradigms in the inter- and intra-domain CDR tasks.

## 2 RELATED WORK

### 2.1 Cross Domain Recommendation

Cross-domain recommendation (CDR) aims to enhance the performance of the target domain by utilizing data from other related domains. Recent surveys [31, 36] divide CDR tasks into inter- and intra-domain recommendations. For **intra-domain** recommendation, the recommended items are in the same domain as the user, while for **inter-domain** recommendation, the recommended items are from a different domain. Specifically, the inter-domain CDR is also referred to as **cold-start user** recommendation, because the target users do not have historical interactions in the domain.

For the intra-domain task, the pioneering study CMF [24] is naturally based on a unified interaction matrix for both domains and trains one matrix factorization model for both domains. With the development of deep learning, various designs of shared deep layers between domains are developed in the **unified distribution paradigm**, including Conet [9], DDTCTR [16], and BiTGCF [19]. In addition, the attention mechanism is commonly used in the adapters to generate domain-specific representations from trainable unified user embeddings [2, 6, 10, 14]. The **embedding and mapping paradigm** is also adopted in the intra-domain task. For example, DTCDR [34] and its variations [35, 37] combine the representations for overlapped users in both domains for prediction. AFT [6] learns the feature mappings across domains under a generative adversarial network. And CUT [15] utilizes user-similarity constraints to filter useful information for the cross-domain mapping.

For the inter-domain task which focuses on cold-start users, the embedding & mapping paradigm is widely adopted. The model first learns the latent factor representation for each domain, and then a mapping function is trained to establish the relationships between the latent space of domains, which is further used to map users to the target domain for cold-start recommendations, such as EMCDDR [21], SSCDDR [11], and PTUPCDR [38]. They rely on MLPs or linear mapping, combined with other user-oriented objectives. Recently, Diffusion models [27] are also adopted as the mapping function for inter-domain recommendations.

Among these CDR frameworks, UniCDR [2] explores the initial attempts to tackle multiple CDR tasks at the same time. Also, by utilizing domain-masking and contrastive loss, UniCDR encourages the domain-shared aggregator to extract the domain-invariant information. However, this additional constraint can interfere with the best recommendations for both domains. Similarly, UCLR [29] also maintains a global user embedding using a contrastive auto-encoder and adopts LoRA as the domain-specific adapter.

Inspired by these previous researches, we follow the unified distribution paradigm to tackle the inter- and intra-domain recommendation tasks concurrently. Unlike existing methods that rely on traditional neural architectures, we leverage conditional diffusion models. Diffusion models share the same target between distribution generation and recommendation process, addressing the

challenge of conflicts between unified and domain-specific objectives in the unified distribution paradigm in CDR. Our innovations also lie in designing effective domain-related condition generators, which fully exploit the diffusion model's capability to model complex distributions for effective CDR.

## 2.2 Diffusion Models for Recommendation

The diffusion probabilistic models (DPMs) [7] provide strong theoretical guarantees for modeling complex data distributions through iterative refinement of a noise distribution. It has achieved remarkable success in various domains such as computer vision and natural language processing [17, 26]. Recently, researchers have applied diffusion models to recommendation systems in various ways. DiffRec [25] utilizes DPMs to model the one-hot user embedding distribution. It corrupts users' historical interactions in forward steps and tries to recover the original interactions. DreamRec [30] reshapes sequential recommendation as a learning-to-generate paradigm using a conditional diffusion model, where the condition is the history sequence, and the generative goal is the target item. Here the diffusion model depicts the distribution of item embeddings. DPMs are also applied to model the general distribution of interaction sequences from the user for sequential recommendations [5, 18]. Conditional diffusion models are also applied to the sequential recommendation scenario [26], where preceding sequence representations from another encoder are adopted as the conditions. A recent model, DiffCDR [27], applies the diffusion model to CDR tasks. However, it follows the embedding and mapping paradigm, and DPMs only work as a more adaptive mapping function, overlooking the diffusion models' powerful capability in modeling complex distributions.

Our CD-CDR aims to achieve a unified item distribution across domains using conditional DPMs, which mitigates the noise problem and the cold start problem in both inter- and intra-domains. In this way, we fully exploit and adapt the power of diffusion models for cross-domain recommendations.

## 3 TASK DEFINITIONS

Let  $\mathcal{U}^A$  and  $\mathcal{U}^B$  denote the user sets in source domain A and target domain B respectively, and  $\mathcal{I}^A$  and  $\mathcal{I}^B$  denote the corresponding item sets. For each user  $u \in \mathcal{U}^A$  (or  $\mathcal{U}^B$ ), we have their interaction history  $\mathcal{H}_u^A = \{i_1, i_2, \dots, i_n\}$  where  $i_k \in \mathcal{I}^A$  (or  $\mathcal{I}^B$ ).

In cross-domain recommendations, two types of recommendation tasks are considered:

- **Intra-domain recommendation:** For target user  $u \in \mathcal{U}^B$ , recommend target items  $i \in \mathcal{I}^B$ .
- **Inter-domain recommendation:** For source user  $u \in \mathcal{U}^A$  with no historical interactions in domain B, recommend target items  $i \in \mathcal{I}^B$ .

Both tasks are commonly adopted in the CDR scenarios. The intra-domain task evaluates the cross-domain recommenders' ability to enhance the target domain recommendations with information from other domains, while the inter-domain recommendations are essential for expanding new users and effective transfer of user preferences across domains. The inter-domain recommendation task is particularly challenging as it involves *cold-start* users who have no interaction history in the target domain.

## 4 MODEL FRAMEWORK

### 4.1 Model Overview

To address inter- and intra-domain CDR tasks with the unified distribution paradigm, we develop **CD-CDR**, an innovative method that adopts a conditional diffusion model to create a unified item distribution across different domains. The overall framework of CD-CDR is shown in Figure 2.

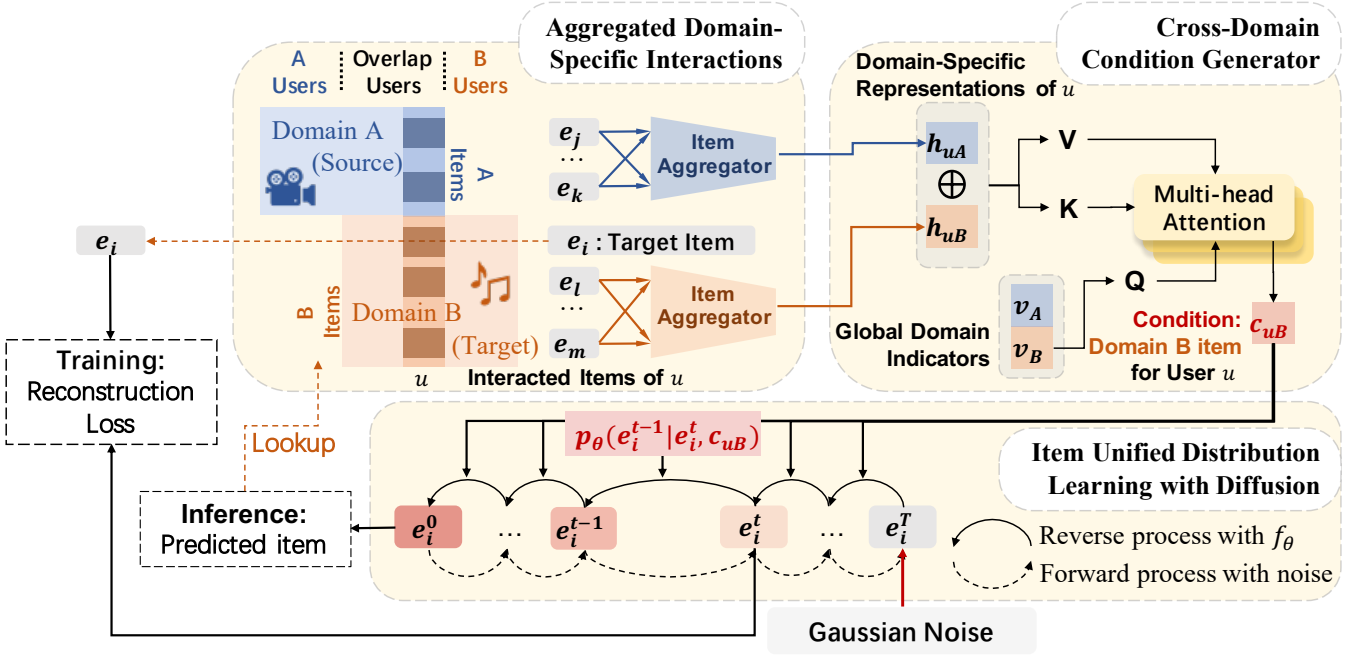
The core component of CD-CDR is a **conditional diffusion model**, which is adopted to learn a unified item distribution across domains. The diffusion model works as the bridge between different domains, providing a comprehensive understanding of item representations. Centered around the diffusion model, CD-CDR consists of three key modules: domain-specific interaction aggregation, cross-domain condition generator, and item-unified distribution learning with the diffusion model. The first two modules prepare the conditions for the diffusion model, ensuring the conditions express both domain-specific preferences and cross-domain shifts. In this way, CD-CDR can handle both inter-domain and intra-domain recommendation tasks. The unified distribution is learned with the last module, which directly involves the diffusion model.

To be specific, firstly, the domain-specific interaction aggregation module takes user interaction histories  $\mathcal{H}_u^A$  and  $\mathcal{H}_u^B$  as input and outputs aggregated user representations  $\mathbf{h}_u^A \in \mathbb{R}^d$  and  $\mathbf{h}_u^B \in \mathbb{R}^d$  for each domain, as shown in Section 4.2. Then, the cross-domain condition generator, as detailed in Section 4.3, takes aggregated representations  $\mathbf{h}_u^A, \mathbf{h}_u^B$ , as well as domain indicators  $\mathbf{v}_A/\mathbf{v}_B$  as input. The output of the generator is domain-aware conditions  $\mathbf{c}_{uB}$  or  $\mathbf{c}_{uA}$ . Finally, the item-unified distribution learning module takes condition  $\mathbf{c}_{uA}$  (or  $\mathbf{c}_{uB}$ ) and Gaussian noise  $\mathbf{e}_t^T$  as input, and generates the denoised item embedding through the diffusion process defined in Section 4.4, providing recommendations aware of information from both domains.

During model training, the final recommendations are generated by applying these modules sequentially, with a reconstruction loss (Equation 6) supervising the learning process. In the inference phase, CD-CDR utilizes the diffused item representations to make accurate and relevant recommendations by grounding to the target domain items, which we will discuss in Section 4.5. To summarize, CD-CDR leverages the strengths of diffusion models, such as their ability to handle noise and filter out irrelevant domain-specific information, as well as their capacity to maintain consistent item distributions. As a result, by incorporating historical interactions from the source domain, our model can accurately recommend items for active and new users in the target domain.

### 4.2 Domain-Specific Interaction Aggregation

When learning unified item distributions across domains, it is crucial to first capture the unique characteristics and behavioral patterns within each domain for users. This is because while we aim for a unified distribution, the user preferences and interaction patterns can vary significantly between domains. By independently aggregating interactions from each domain before combining them, we ensure that domain-specific nuances are preserved and can be effectively incorporated into the unified distribution through the conditional diffusion process. This design choice aligns with our goal of maintaining domain awareness while learning shared



**Figure 2: The CD-CDR model generates unified item distributions across domains with a conditional diffusion model. It includes three key modules: Domain-specific interaction aggregation, Cross-domain condition generator, and Item-unified distribution learning with diffusion.**

knowledge, without requiring extra supervision signals to balance domain-specific and unified objectives.

In the domain-specific interaction aggregation module, we independently aggregate interactions from each domain, thereby enabling the model to capture domain-specific nuances effectively. To illustrate, interactions from domain A and domain B undergo separate aggregation processes, ensuring that the resulting representations encapsulate the unique characteristics of each domain. This approach enables CD-CDR to generate personalized recommendations that cater to the diverse needs of users across different domains. For each domain  $d \in \{A, B\}$ , we aggregate user interactions using a self-attention mechanism. Given a user  $u$ 's interaction history  $\mathcal{H}_u^d = \{i_1, i_2, \dots, i_n\}$  where  $i_k \in \mathcal{I}^d$ , we compute the attention weights  $\alpha_{ij}$  between items  $i$  and  $j$ :

$$\alpha_{ij} = \text{softmax} \left( \frac{(\mathbf{W}_Q \mathbf{e}_i)(\mathbf{W}_K \mathbf{e}_j)^T}{\sqrt{d_k}} \right) \quad (1)$$

where  $\mathbf{W}_Q, \mathbf{W}_K$  are learnable projection matrices and  $d_k$  is the dimension of the key vectors. The final aggregated representation for user  $u$  in domain  $d$  is:

$$\mathbf{h}_u^d = \sum_{i,j \in \mathcal{H}_u^d} \alpha_{ij} (\mathbf{W}_V \mathbf{e}_j) \quad (2)$$

where  $\mathbf{W}_V$  is the value projection matrix.

In contrast to existing approaches such as DreamRec, CD-CDR distinguishes itself through its explicit consideration of cross-domain

dynamics and the incorporation of diffusion-based item distribution learning. By leveraging the strengths of diffusion and domain-specific aggregation, CD-CDR offers a unique solution to the challenges of cross-domain recommendation, enabling more accurate and personalized recommendations tailored to the preferences of users across diverse domains.

### 4.3 Cross-Domain Condition Generator

A key challenge in cross-domain recommendation is effectively combining information from different domains while maintaining domain awareness. Traditional approaches often require extra supervision signals to balance unified and domain-specific objectives. Our cross-domain condition generator addresses this challenge by leveraging the natural conditioning mechanism of diffusion models. By using domain indicators and attention mechanisms, we can guide the diffusion process to generate recommendations that are both domain-aware and benefit from cross-domain knowledge, without requiring additional supervision signals or explicit domain adapters. Depending on the domain  $d \in \{A, B\}$  of the target item  $i$ , we have a trainable global domain vector  $\mathbf{v}_d$  as the indicator. It is worth noting that while target items during training can originate from either the source or target domains, during inference we exclusively generate recommendations for items in the target domain. Given the aggregated user representations  $\mathbf{h}_u^A$  and  $\mathbf{h}_u^B$  from both domains, we compute the condition of the diffusion process through a domain-specific attention mechanism:

$$\mathbf{c}_{ud} = \text{Attention}(\mathbf{v}_d, [\mathbf{h}_u^A; \mathbf{h}_u^B]) \quad (3)$$



where  $[\cdot; \cdot]$  denotes concatenation and the attention weights are computed as:

$$\text{Attention}(\mathbf{q}, \mathbf{k}) = \text{softmax}\left(\frac{\mathbf{q}\mathbf{k}^T}{\sqrt{d}}\right) \mathbf{k} \quad (4)$$

Here,  $\mathbf{v}_d$  serves as the query vector and the concatenated user representations serve as both key and value vectors. This mechanism allows the model to dynamically adjust the weights assigned to each domain's information based on their relevance to the target domain.

The generated condition vector  $\mathbf{c}_{ud}$  encapsulates both the user's preferences and domain-specific characteristics, which is then used to guide the diffusion process in generating recommendations. By incorporating both domain indicators and user-specific attention, the model can effectively balance information from different domains while maintaining domain awareness in the recommendation process.

#### 4.4 Item Unified Distribution Learning with Diffusion

The core innovation of our approach lies in using diffusion models to learn a unified item distribution across domains. Unlike traditional unified distribution approaches that require extra supervision signals to maintain consistency across domains, diffusion models naturally excel at modeling complex distributions and can incorporate domain-specific information through conditioning. This allows us to simultaneously capture shared knowledge between domains while preserving domain-specific characteristics, fundamentally resolving the trade-off between unified and domain-specific objectives that exist in previous approaches. During training, for each piece of data log containing a target item  $i$  in domain  $d$ , we train the diffusion model to generate this target item based on the user's historical interactions from both domains and the domain indicator  $d$ , which is encoded in the condition  $\mathbf{c}_{ud}$ . Importantly, the target items can come from either the source or target domains, enabling the diffusion model to learn a unified distribution that encompasses items from both domains.

Given a target item embedding  $\mathbf{e}_i$  from domain  $d$ , the training process begins by sampling a diffusion time step  $t$  from a uniform distribution over  $[1, \dots, T]$ . Next, we sample Gaussian noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and add this noise to the target embedding:

$$\mathbf{e}_i^t = \sqrt{\alpha_t} \mathbf{e}_i + \sqrt{1 - \alpha_t} \epsilon \quad (5)$$

Then, we sample an indicator with probability  $p_u$  to determine whether to use the condition - with probability  $p_u$  we remove the condition by setting it to  $\Phi$ , and with probability  $1 - p_u$  we keep the condition  $\mathbf{c}_{ud}$ . Finally, we compute the reconstruction loss and perform a gradient update:

$$L_{recon} = \|\mathbf{e}_i - f_\theta(\mathbf{e}_i^t, c, t)\|^2 \quad (6)$$

where  $f_\theta$  is the model's prediction function parameterized by  $\theta$ , and  $c$  is either  $\mathbf{c}_{ud}$  or  $\Phi$  based on the sampled indicator.

The noise schedule is controlled by parameters  $\alpha_t$  and  $\bar{\alpha}_t$ , defined as:

$$\alpha_t = 1 - \beta_t, \bar{\alpha}_t = \prod_{s=1}^t \alpha_s \quad (7)$$

where  $\alpha_s$  are predefined hyperparameters. The reweighted noise scale factor  $\tilde{\beta}_t$  is:

$$\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \quad (8)$$

#### 4.5 Item Generation in the Inference Phase

The inference phase is critical for demonstrating the practical value of our unified distribution approach, particularly for cold-start users in inter-domain recommendation scenarios. While we train the model on both domains to learn comprehensive item distributions, during inference we focus exclusively on generating recommendations in the target domain. This design choice allows us to leverage the rich knowledge captured in the unified distribution while ensuring that recommendations are properly adapted to the target domain's characteristics, addressing both intra-domain and inter-domain recommendation tasks effectively. During inference, we utilize the user's historical interactions from both domains, while fixing the target domain indicator  $d = B$ . The model leverages the learned unified distribution to generate potential items of interest specifically in the target domain, as we evaluate the recommendation performance only on the target domain.

The inference process begins by sampling initial noise  $\mathbf{e}_i^T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Then, for each time step  $t$  from  $T$  down to 1, we apply the reverse diffusion process in two stages. First, we compute the prediction using both conditional and unconditional models:

$$\tilde{f}_\theta(\mathbf{e}_i^t, \mathbf{c}_{ud}, t) = (1 + w)f_\theta(\mathbf{e}_i^t, \mathbf{c}_{ud}, t) - wf_\theta(\mathbf{e}_i^t, \Phi, t)$$

where  $w$  is a hyperparameter controlling the balance between conditional and unconditional predictions. Next, we update the item embedding using:

$$\mathbf{e}_i^{t-1} = \frac{\sqrt{\bar{\alpha}_{t-1}}\tilde{\beta}_t}{1 - \bar{\alpha}_t} \tilde{f}_\theta(\mathbf{e}_i^t, \mathbf{c}_{ud}, t) + \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{e}_i^t + \sqrt{\tilde{\beta}_t} \mathbf{z} \quad (9)$$

where  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  is a random noise vector.

The final denoised embedding  $\mathbf{e}_i^0$  represents the generated item recommendation in the target domain. This process leverages the unified item distribution learned during training to generate recommendations that are consistent with the target domain while being conditioned on user preferences through  $\mathbf{c}_{ud}$ . By learning from both domains during training but focusing on the target domain during inference, CD-CDR effectively addresses the challenges of cross-domain recommendation, enabling high-quality recommendations even for users with no interaction history in the target domain.

## 5 EXPERIMENTS

In this section, we validate the proposed CD-CDR model on two real-world datasets and two CDR tasks: inter-domain and intra-domain (i.e., cold-start user) recommendations. In both tasks, CD-CDR is compared against state-of-the-art single and cross-domain recommenders, including both embedding & mapping and unifying & adapting paradigms.

Moreover, we compare the learned unified item distributions of CD-CDR with other baselines. We also indicate the effectiveness of each module in CD-CDR and its sensitivity to the hyper-parameters.

**Table 1: Dataset statistics. Each pair of domains indicates a cross-domain recommendation task. The subscript  $o$  indicates overlap.**

Dataset	Domain	$ \mathcal{U} $	$ \mathcal{I} $	# Clicks	$ \mathcal{U}_o $	$ \mathcal{I}_o $
Amazon	Sports	35,599	18,358	296,337	3,908	704
	Cloth	39,388	23,034	278,677		
	Video	24,034	10,673	231,780	999	0
Douban	Cloth	39,388	23,034	278,677		
	Music	16,041	40,405	1,140,090	14,000	0
	Movie	22,254	27,432	2,760,500		

## 5.1 Experimental Settings

**5.1.1 Datasets & Evaluation Metrics.** We evaluate our model on two real-world cross-domain datasets, with three domain pairs in total. For each pair, we consider both domains as the target domain alternatively, resulting in six cross-domain recommendation tasks. The datasets include:

- **Amazon**<sup>1</sup>: A large-scale e-commerce dataset with user interactions from multiple domains. We choose two domain pairs, Cloth&Sports, and Cloth&Video, and conduct cross-domain recommendations respectively, where Cloth and Sports are more closely related while Cloth and Video share less common knowledge.
- **Douban**<sup>2</sup>: Douban is a music and movie online platform, where we consider two cross-domain tasks with music and movie as target/source domains, respectively.

Both datasets are widely used for cross-domain recommendations [2, 3, 15, 19]. As shown in the dataset statistics in Table 1, all domain pairs share overlapping users ( $|\mathcal{U}_o|$ ). Following the previous work [15], we transform the ratings into implicit data where each entry is marked as 0 or 1 according to whether the user has interacted with the item. We filter the dataset to keep users and items with at least 5 interactions. The source domain is split by 8:2 for training and validation for fair comparisons with previous CDR methods with source phases, such as UCLR [29] and CUT [15]. As for the target domain, intra-domain recommendations are conducted with the user history split by the ratio of 8:1:1 for training, validation, and testing for each user. As for the inter-domain task, we follow the previous setting [2, 4] where the test set consists of 10% of the overlapping users, whose target domain interactions are saved for testing.

For evaluation, a full ranking setting is utilized, where the recommendation is conducted on all items in the datasets. We evaluate all CDR tasks by HR@10 and NDCG@10 on the target domain; both are commonly used evaluation metrics.

**5.1.2 Compared Baselines.** We compare CD-CDR with various single-domain and cross-domain baselines. Single-domain baselines are compared in the inter-domain CDR task. They are trained solely on the target dataset, including the classical MF recommender and SOTA diffusion-based models:

- **MF** [13] is the classic matrix factorization model that represents users and items with latent factors.

<sup>1</sup>[http://jmcauley.ucsd.edu/data/amazon/index\\_2014.html](http://jmcauley.ucsd.edu/data/amazon/index_2014.html)

<sup>2</sup><https://recbole.s3-accelerate.amazonaws.com/CrossDomain/Douban.zip>

- **DiffRec** [25] corrupts users' one-hot embedding in forward steps and tries to recover the original interactions. The recovered user embedding is then used to calculate dot products with candidate items.
  - **DreamRec** [30] is one of the SOTA diffusion-based recommenders, which utilizes history interaction item sequence as the condition and generates an oracle item from noise prior.
- Note that we mainly compare with diffusion-based single-domain methods since our target is to validate the effectiveness of applying diffusion models for CDR tasks rather than beating sophisticated single-domain models.

For cross-domain baselines, we compare our CD-CDR framework with classical and SOTA algorithms in both inter-domain and intra-domain CDR tasks. Following previous definitions, the baselines can be classified into the embedding&mapping paradigm and the unified distribution paradigm. The baselines with the unified distribution paradigm include:

- **CMF** [24] is a widely-used classical cross-domain recommender. It is MF trained on both domains with different weights in prediction loss.
- **UniCDR** [2] transfers the most relevant domain-shared information across domains by its domain-shared and specific user embeddings and encourages the information transfer using interaction-level contrastive learning.
- **UCLR** [29] leverages a contrastive dual-stream collaborative autoencoder to provide consistent embeddings for inactive users in all domains, thus enhancing the robustness of CDR. It also employs Low-Rank Adaptation (LoRA) to improve the performance in each specific domain.

For embedding&mapping paradigms, we compare with:

- **CUT** [15] induces explicit constraints based on target-domain user similarities to mitigate negative transfer. CUT is a framework, for the backbone we adopt MF and LightGCN following the original authors.
- **DiffCDR** [28] utilize diffusion models as the mapping function between user representations across domains, with MF as backbone for each domain.

**5.1.3 Implementation Details.** We implement all models using the RecBoleCDR [33] library. User history is set to 10 interactions, and the embedding size is fixed at 64 for all models. We carefully tune hyper-parameters on all baselines based on the performance of the validation set. In our proposed CD-CDR, the batch size is 2048, the learning rate of Adam [12] optimizer is set as 0.001, the diffusion step is tuned in [1, 10, 100, 1000], the training phase unconditional sample probability  $p$  is 0.1 and the inference phase unconditional weight factor  $w$  is 2. The weight decay is set to 1e-6 for the Amazon dataset and 1e-7 for Douban after grid search, and the BPR loss [22] is adopted for both datasets. The open-sourced repository is released<sup>3</sup>.

## 5.2 Performance on intra-domain CDR

In the intra-domain CDR task, which is the most common task for CDR [2, 15, 24], both users and items are from the target domain. Cross-domain models attempt to leverage auxiliary information

<sup>3</sup><https://anonymous.4open.science/r/CD-CDR-1323>

**Table 2: Performance comparisons on the intra-domain CDR for six cross-domain tasks. All experiments are repeated five times, and \* shows statistical significance (paired t-test with p-value < 0.05) compared with the best baselines. The best performances are in bold, and the second-best results are underlined.**

Dataset	Domain: Source → Target	Metrics (@10)	Single Domain Methods			Cross Domain Methods						Our Method
			MF	DiffRec	DreamRec	CMF	UniCDR	UCLR	CUT-MF	CUT-LightGCN	DiffCDR	
Amazon	Cloth → Sports	Recall	0.0492	0.0613	0.0602	0.0545	0.0624	0.0638	0.0601	<u>0.0653</u>	0.0573	<b>0.0695*</b>
		NDCG	0.0270	0.0312	0.0307	0.0293	0.0340	0.0351	0.0335	<u>0.0364</u>	0.0302	<b>0.0384*</b>
	Sports → Cloth	Recall	0.0243	0.0357	0.0377	0.0291	0.0433	0.0426	0.0393	<u>0.0441</u>	0.0298	<b>0.0464*</b>
		NDCG	0.0137	0.0198	0.0213	0.0157	0.0239	0.0235	0.0222	<u>0.0252</u>	0.0161	<b>0.0286*</b>
	Cloth → Video	Recall	0.1153	0.1273	0.1301	0.1194	0.1249	0.1282	0.1275	<u>0.1303</u>	0.1027	<b>0.1341*</b>
		NDCG	0.0623	0.0689	0.0692	0.0644	0.0684	0.0693	0.0704	<u>0.0720</u>	0.0625	<b>0.0752*</b>
	Video → Cloth	Recall	0.0243	0.0357	0.0377	0.0246	0.0349	0.0374	0.0362	<u>0.0381</u>	0.0253	<b>0.042*</b>
		NDCG	0.0137	0.0198	<u>0.0213</u>	0.0136	0.0191	0.0210	0.0203	<u>0.0213</u>	0.0143	<b>0.0226*</b>
Douban	Movie → Music	Recall	0.1004	0.1176	<b>0.1275</b>	0.0944	0.1073	0.1189	0.1238	0.1205	0.0935	0.1244
		NDCG	0.0733	0.0933	<u>0.0962</u>	0.0725	0.0754	0.0927	0.0952	0.0946	0.0719	<b>0.0981</b>
	Music → Movie	Recall	0.1053	0.1278	0.1306	0.0946	0.1095	0.1231	0.1390	<u>0.1393</u>	0.0929	<b>0.1418</b>
		NDCG	0.0997	<b>0.1469</b>	0.1408	0.1031	0.0994	0.1344	0.1413	0.1437	0.1017	<u>0.1462</u>

**Table 3: Performance comparisons on the inter-domain cold-start CDR for six cross-domain tasks. All the notations are the same as Table 2.**

Dataset	Domain: Source → Target	Metrics (@10)	Cross Domain Methods						Our Method
			CMF	UniCDR	UCLR	CUT-MF	CUT-LightGCN	DiffCDR	
Amazon	Cloth → Sports	Recall	0.0954	0.0523	0.0619	0.0975	0.0407	0.1045	<b>0.1172*</b>
		NDCG	0.1341	0.0550	0.0734	0.1082	0.0427	0.1436	<b>0.1470</b>
	Sports → Cloth	Recall	0.0954	0.0655	0.0892	0.0863	0.0149	<u>0.1012</u>	<b>0.1032</b>
		NDCG	0.1341	0.0812	0.1283	0.1216	0.0157	<u>0.1468</u>	<b>0.1535*</b>
	Cloth → Video	Recall	0.0234	0.0230	0.0205	0.0192	<u>0.0244</u>	0.0238	<b>0.0265*</b>
		NDCG	0.0347	0.0346	0.0282	0.0205	<u>0.0372</u>	0.0352	<b>0.0376</b>
	Video → Cloth	Recall	0.0084	<u>0.0157</u>	0.0136	0.0072	0.0138	0.0113	<b>0.0191*</b>
		NDCG	0.0102	<b>0.0206</b>	0.0181	0.0089	0.0150	0.0134	0.0198
Douban	Movie → Music	Recall	0.0328	0.0319	<u>0.0336</u>	0.0322	0.0317	0.0301	<b>0.0342*</b>
		NDCG	<u>0.1345</u>	0.1278	0.1324	0.1303	0.1289	0.1264	<b>0.1398*</b>
	Music → Movie	Recall	<u>0.0489</u>	0.0466	0.0485	0.0487	0.0432	0.0476	<b>0.0513*</b>
		NDCG	<u>0.3070</u>	0.2962	0.2978	0.3034	0.2762	0.3052	<b>0.3300*</b>

from the source domain to enhance target-domain recommendation accuracy. Meanwhile, single-domain methods can be applied in this scenario due to the availability of history interactions for users. We compare our method against both single-domain and cross-domain baselines, as shown in Table 2.

The results on six intra-domain CDR tasks indicate that our proposed CD-CDR has achieved significant improvements against the best baselines in most scenarios, especially on four tasks in the Amazon dataset. For the Amazon datasets, the CUT framework with the LightGCN backbone is the best baseline due to its explicit constraint to alleviate negative transfer when transferring source-domain information. Our method exceeds CUT-LightGCN significantly by 5.49%, 13.4%, 4.44%, and 6.10% in terms of NDCG@10 for four tasks in the Amazon dataset, respectively. This advantage of CD-CDR reveals the effectiveness of using the diffusion model with high representational power to encode items from two domains into a unified distribution. Meanwhile, DreamRec, which also employs conditional diffusion models, demonstrates impressive

performance, highlighting the effectiveness of diffusing conditions. However, DiffCDR, which utilizes the diffusion model as the mapping function, fails to achieve good performances. It shows the limitations of adopting diffusion models solely as the mapping function in CDR. In CD-CDR, the unified distribution paradigm is adopted with specifically designed conditions for cross-domain diffusion models, resulting in the best performance.

The Douban datasets are denser, with more interactions in the target domains. Therefore, the impact of the source domain data is relatively limited for Douban. Thus, the information from source domains introduced by cross-domain models fails to further improve target domain performance, and single-domain diffusion-based recommenders generally achieve superior performances. Our proposed CD-CDR, while not consistently surpassing single-domain methods, outperforms all cross-domain baselines, benefiting from its powerful representation ability for both domains.




Figure 3: Ablation study of the Domain-Specific Aggregator.

### 5.3 Performance on inter-domain CDR

Inter-domain CDR aims to recommend source-domain items to target-domain users, presenting a **cold-start** scenario due to the absence of historical interactions [31]. This task serves as a critical testbed for evaluating cross-domain recommenders' transfer learning capabilities, while single-domain methods are inherently inapplicable. Table 3 presents the comparison between CD-CDR and all cross-domain baselines. Note that as users can have multiple positive interactions in each candidate list, NDCG may exceed Recall, and the results are incomparable with Table 2.

For this challenging yet valuable CDR task, our proposed CD-CDR significantly outperforms all existing cross-domain approaches across various evaluation metrics on most metrics. These improvements underscore the effectiveness of employing a diffusion model with a high representational capacity to integrate unified item representations. Furthermore, the domain identifiers, working as condition generators for the diffusion process, effectively distinguish between the behaviors observed in the source and target domains. This enables CD-CDR to capture nuanced user preferences across different domains, providing better recommendations for cold users. Comparing different baselines, DiffCDR, CUT-MF, and CMF all have impressive performance on the first and last two tasks. They are all based on the MF backbone, showing the importance of modeling collaborative information in the inter-domain CDR task. Despite this, CD-CDR still achieves the best results, indicating that collaborative information can be effectively learned by the unified item distributions and history-aggregated conditions in CD-CDR.

Looking into the results for intra- and inter-domain tasks, we find that neither the embedding & mapping paradigm nor the unified distribution paradigm consistently outperforms the other. Our method, CD-CDR, adopts the unified distribution paradigm while addressing its main limitation - the trade-off between domain-shared and domain-specific objectives - through a conditional diffusion model. Thus, CD-CDR achieves strong performance on both tasks.

### 5.4 Ablation Study

In this ablation study, we focus on the effectiveness of the condition generation process in CD-CDR, including the domain-specific




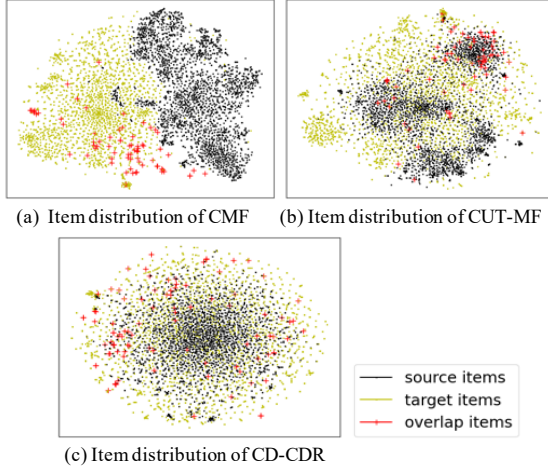
Figure 4: Ablation study of the Cross-Domain Condition Generator. w/o is short for “without”.

aggregation module and the cross-domain condition generation module. All ablation studies are conducted on the Cloth ↔ Sports domains in Amazon for both inter- and intra-domain CDR tasks.

For the domain-specific aggregation module, we compare with the variant without modeling the domain-specific user history. To be specific, the variant mixes historical interactions from both domains into a sequence to generate one representation for each user, which is called *Directly Mix Domain Data*. The results in Figure 3 show that CD-CDR consistently outperforms the *Directly Mix Domain Data* variant across all evaluation metrics, indicating the domain-specific interaction aggregation module plays a crucial role in enhancing the recommendation performance of CD-CDR. By independently aggregating interactions from each domain, the aggregation module ensures that the output representations capture the unique features of each domain, thus enhancing the model's ability to generate domain-dependent personalized recommendations. Directly mixing interactions from both domains compromises the model's ability to capture domain-specific user preferences, thus degrading its overall performance.

Ablation study on the cross-domain condition generator is conducted with two variants: *without Domain-specific User Attention* and *without Domain Indicator*. The domain-specific representations are directly concatenated for the variant without attention, and the global indicators are replaced with an all-one vector for the variant without indicators. Figure 4 shows CD-CDR outperforms its two variants in almost all tasks and metrics, demonstrating the importance of our design of attention-based domain indicator condition generation. Specifically, condition generation without attention leads to a large performance decrease, indicating the importance of constructing a condition with distinguishable preferences for different domains. On the other hand, removing domain indicators has a limited impact on the performance. This is because the unified distribution is shared across domains, and the generated representations from the unified diffusion model will be grounded to the target domain items in the inference stage. Therefore, CD-CDR still works under the unified distribution even without the indicator, illustrating the power of the diffusion model for distribution generation from another perspective.





**Figure 5: Visualization of item distributions from source, target, and overlap items on CMF, CUT-MF, and CD-CDR.**

### 5.5 Domain Item Distribution Visualization


To illustrate the unified item distribution in CD-CDR, we visualize the item embedding distributions for different domains with three methods: CMF [24], CUT-MF [15], and our proposed CD-CDR. The visualization aims to provide a clear comparison of how these methods handle item embeddings across domains.

In Figure 5, we plot the item embeddings for the Amazon Cloth  $\rightarrow$  Sports dataset using t-SNE. Both CMF and CUT-MF show dispersed item embedding distributions from different domains, especially for the source domain items. In CMF, items from different domains are almost completely separated, with overlapping items primarily in the target domain. As a result, items from different domains are treated differently, leading to suboptimal recommendations, especially for cold start users who have no interactions in the target domain. In CUT-MF, since a mapping function is adopted to transfer items across domains, the item distributions are overlapped. However, the source item distribution is dispersed and clustered, which may lead to challenges in providing accurate recommendations for users who interact with items from multiple domains. In contrast, our proposed CD-CDR method demonstrates a more unified and coherent distribution of item embeddings across both domains. The item embeddings from source and target domains are more closely aligned, indicating that the diffusion process successfully learns a unified representation of item distributions. This alignment ensures that items from different domains are represented consistently and coherently, which is crucial for improving the recommendation performance, especially for cold start users.

### 5.6 Parameter Sensitivity

To ensure the robustness and effectiveness of our proposed CD-CDR model, we conducted a parameter sensitivity analysis on the Cloth  $\rightarrow$  Sports CDR task of the Amazon dataset. This analysis focuses on two key parameters: the Unconditional Weight Factor and the Diffusion Step. The results are shown in Figure 6.

The Unconditional Weight Factor controls the influence of unconditional (i.e., domain-agnostic) information in the recommendation



**Figure 6: Parameter sensitivity of the Unconditional Weight Factor and the Diffusion Step.**

process. The best performance was observed when the Unconditional Weight Factor was set to 2, and both larger and smaller weights led to slightly worse results. However, the differences in performance between the values were relatively small, indicating that the model is not highly sensitive to the unconditional weight parameter. The Diffusion Step determines the number of steps in the diffusion process, which is critical for the model’s ability to capture complex interactions and patterns across domains. We evaluated the model’s performance with different values of the Diffusion Step. The best performance was achieved when the Diffusion Step was set to 100. A larger Diffusion Step resulted in significantly worse performance, likely due to over-diffusion leading to instability and noise in the recommendations. A Diffusion Step of 100 allows the model to capture a sufficient amount of interaction complexity without introducing excessive noise or instability. The poor performance at 1000 steps also aligns with the findings in DreamRec [30], where too many diffusion steps can lead to overfitting and unstable recommendations.

## 6 CONCLUSION

In this paper, we propose a new model, CD-CDR, for cross-domain recommendation (CDR) tasks. To the best of our knowledge, it is the first attempt to explore the unified distribution paradigm in CDR scenarios with the conditional diffusion model. CD-CDR generates conditions for diffusion models based on domain-specific preferences and cross-domain information, sharing the same target between diffusion and recommendation processes. Thus, our proposed approach effectively addresses the long-standing challenge of conflicts between unified and domain-specific objectives in the unified distribution paradigm in the CDR scenario. Consequently, CD-CDR applies to both inter-domain and intra-domain (i.e., cold-start) CDR tasks.

Extensive experiments are conducted on both inter- and intra-domain tasks with six cross-domain pairs from two real-world datasets, where our proposed CD-CDR is compared with both state-of-the-art single and cross-domain baselines. The encouraging results of CD-CDR overall task settings and datasets highlight its general efficacy in improving recommendation accuracy in cross-domain recommendation scenarios. Overall, this study contributes to advancing the understanding and application of generative models in CDR systems as distribution generators, paving the way for more effective and personalized cross-domain recommendations for users.

## REFERENCES

- [1] 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*. ACM, Atlanta GA USA, 138–147. <https://doi.org/10.1145/3511808.3557262>
- [2] Jiangxia Cao, Shaoshuai Li, Bowen Yu, Xiaobo Guo, Tingwen Liu, and Bin Wang. 2023. Towards Universal Cross-Domain Recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. ACM, Singapore Singapore, 78–86. <https://doi.org/10.1145/3539597.3570366>
- [3] 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*. ACM, Madrid Spain, 267–277. <https://doi.org/10.1145/3477495.3531967>
- [4] Jiangxia Cao, Jiawei Sheng, Xin Cong, Tingwen Liu, and Bin Wang. 2022. Cross-Domain Recommendation to Cold-Start Users via Variational Information Bottleneck. 2022 *IEEE 38th International Conference on Data Engineering (ICDE)* (2022). <https://doi.org/10.48550/arXiv.2203.16863>
- [5] Hanwen Du, Huanhuan Yuan, Zhen Huang, Pengpeng Zhao, and Xiaofang Zhou. 2023. Sequential Recommendation with Diffusion Models. <http://arxiv.org/abs/2304.04541> arXiv:2304.04541 [cs].
- [6] Xiaobo Hao, Yudan Liu, Ruobing Xie, Kaikai Ge, Linyao Tang, Xu Zhang, and Leyu Lin. 2021. Adversarial Feature Translation for Multi-domain Recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM, Virtual Event Singapore, 2964–2973. <https://doi.org/10.1145/3447548.3467176>
- [7] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. <http://arxiv.org/abs/2006.11239> arXiv:2006.11239 [cs, stat].
- [8] Jonathan Ho and Tim Salimans. 2022. Classifier-Free Diffusion Guidance. <http://arxiv.org/abs/2207.12598> arXiv:2207.12598 [cs].
- [9] 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*. ACM, Torino Italy, 667–676. <https://doi.org/10.1145/3269206.3271684>
- [10] Yuchen Jiang, Qi Li, Han Zhu, Jinbei Yu, Jin Li, Ziru Xu, Huihui Dong, and Bo Zheng. 2022. Adaptive Domain Interest Network for Multi-domain Recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. ACM, Atlanta GA USA, 3212–3221. <https://doi.org/10.1145/3511808.3557137>
- [11] SeongKu Kang, Junyoung Hwang, Dongha Lee, and Hwanjo Yu. 2019. Semi-Supervised Learning for Cross-Domain Recommendation to Cold-Start Users. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (2019).
- [12] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014).
- [13] Yehuda Koren, Robert M. Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42 (2009).
- [14] Chenglin Li, Yuanzhen Xie, Chenyun Yu, Bo Hu, Zang li, Guoqiang Shu, Xiaohu Qie, and Di Niu. 2022. One for All, All for One: Learning and Transferring User Embeddings for Cross-Domain Recommendation. <https://doi.org/10.1145/3539597.3570379> arXiv:2211.11964 [cs].
- [15] Hanyu Li, Weizhi Ma, Peijie Sun, Jiayu Li, Cunxiang Yin, Yancheng He, Guoqiang Xu, Min Zhang, and Shaoping Ma. 2024. Aiming at the Target: Filter Collaborative Information for Cross-Domain Recommendation. arXiv:2403.20296 [cs.IR]
- [16] P. Li and Alexander Tuzhilin. 2019. DDTCDR: Deep Dual Transfer Cross Domain Recommendation. *Proceedings of the 13th International Conference on Web Search and Data Mining* (2019).
- [17] Xiang Lita Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. 2022. Diffusion-LM Improves Controllable Text Generation. <http://arxiv.org/abs/2205.14217> arXiv:2205.14217 [cs].
- [18] Zihao Li, Aixin Sun, and Chenliang Li. 2023. DiffuRec: A Diffusion Model for Sequential Recommendation. <http://arxiv.org/abs/2304.00686> arXiv:2304.00686 [cs].
- [19] 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*. ACM, Virtual Event Ireland, 885–894. <https://doi.org/10.1145/3340531.3412012>
- [20] Xihui Liu, Dong Huk Park, Samaneh Azadi, Gong Zhang, Arman Chopikyan, Yuxiao Hu, Humphrey Shi, Anna Rohrbach, and Trevor Darrell. 2022. More Control for Free! Image Synthesis with Semantic Diffusion Guidance. <http://arxiv.org/abs/2112.05744> arXiv:2112.05744 [cs].
- [21] Tong Man, Huawei Shen, Xiaolong Jin, and Xueqi Cheng. 2017. Cross-Domain Recommendation: An Embedding and Mapping Approach. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Melbourne, Australia, 2464–2470. <https://doi.org/10.24963/ijcai.2017/343>
- [22] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. *ArXiv* abs/1205.2618 (2009).
- [23] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. <http://arxiv.org/abs/2112.10752> arXiv:2112.10752 [cs].
- [24] Ajit P. Singh and Geoffrey J. Gordon. 2008. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, Las Vegas Nevada USA, 650–658. <https://doi.org/10.1145/1401890.1401969>
- [25] Wenjie Wang, Yiyan Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. 2023. Diffusion Recommender Model. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, Taipei Taiwan, 832–841. <https://doi.org/10.1145/3539618.3591663>
- [26] Yu Wang, Zhiwei Liu, Liangwei Yang, and Philip S. Yu. 2023. Conditional Denoising Diffusion for Sequential Recommendation. <http://arxiv.org/abs/2304.11433> arXiv:2304.11433 [cs].
- [27] Yuner Xuan. 2024. Diffusion Cross-domain Recommendation. arXiv:2402.02182 [cs.IR]
- [28] Yuner Xuan. 2024. Diffusion Cross-domain Recommendation. <http://arxiv.org/abs/2402.02182> arXiv:2402.02182 [cs].
- [29] Wenhao Yang, Yingchun Jian, Yibo Wang, Shiyin Lu, Lei Shen, Bing Wang, Haihong Tang, and Lijun Zhang. 2024. Not All Embeddings are Created Equal: Towards Robust Cross-domain Recommendation via Contrastive Learning. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*. 3195–3206.
- [30] Zhengyi Yang, Jiancan Wu, Zhicai Wang, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2023. Generate What You Prefer: Reshaping Sequential Recommendation via Guided Diffusion. (2023).
- [31] Tianzi Zang, Yanmin Zhu, Haobing Liu, Ruohan Zhang, and Jiadi Yu. 2022. A Survey on Cross-domain Recommendation: Taxonomies, Methods, and Future Directions. *ACM Transactions on Information Systems* (July 2022), 3548455. <https://doi.org/10.1145/3548455>
- [32] Tong Zhang, Chen Chen, Dan Wang, Jie Guo, and Bin Song. 2023. A VAE-Based User Preference Learning and Transfer Framework for Cross-Domain Recommendation. *IEEE Transactions on Knowledge and Data Engineering* 35, 10 (2023), 10383–10396. <https://doi.org/10.1109/TKDE.2023.3253168>
- [33] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Kaiyuan Li, Yushuo Chen, Yujie Lu, Hui Wang, Changxin Tian, Xingyu Pan, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang, and Ji-Rong Wen. 2021. Recbole: Towards a unified, comprehensive and efficient framework for recommendation algorithms. In *CIKM*.
- [34] Feng Zhu, Chaochao Chen, Yan Wang, Guanfeng Liu, and Xiaolin Zheng. 2019. DTCDR: A Framework for Dual-Target Cross-Domain Recommendation. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (2019).
- [35] Feng Zhu, Yan Wang, Chaochao Chen, Guanfeng Liu, and Xiaolin Zheng. 2020. A Graphical and Attentional Framework for Dual-Target Cross-Domain Recommendation. In *International Joint Conference on Artificial Intelligence*.
- [36] Feng Zhu, Yan Wang, Chaochao Chen, Jun Zhou, Longfei Li, and Guanfeng Liu. 2021. Cross-Domain Recommendation: Challenges, Progress, and Prospects. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Montreal, Canada, 4721–4728. <https://doi.org/10.24963/ijcai.2021/639>
- [37] Feng Zhu, Yan Wang, Jun Zhou, Chaochao Chen, Longfei Li, and Guanfeng Liu. 2021. A Unified Framework for Cross-Domain and Cross-System Recommendations. *IEEE Transactions on Knowledge and Data Engineering* 35 (2021), 1171–1184.
- [38] Yongchun Zhu, Zhenwei Tang, Yudan Liu, Fuzhen Zhuang, Ruobing Xie, Xu Zhang, Leyu Lin, and Qing He. 2021. Personalized Transfer of User Preferences for Cross-domain Recommendation. *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining* (2021).