Dual-Target Disjointed Cross-Domain Recommendation Mediated via Latent User Preferences

Jaime Hieu Do
Singapore Management University
Singapore
dinhhieu.do.2020@smu.edu.sg

Hady W. Lauw
Singapore Management University
Singapore
hadywlauw@smu.edu.sg

ABSTRACT

Users often navigate multiple platforms online, each characterized by its own set of scarce data. Recommender systems face a significant challenge in such fragmented environments. This paper proposes a novel approach to enhance recommendation systems by leveraging connections across distinct yet conceptually similar datasets from multiple platforms. We introduce a unique scenario of dual-target overlapping-free cross-platform recommendation, presenting a bridging mechanism to mutually improve across platforms and learn latent user preferences. Our approach addresses the data sparsity prevalent in each platform and enhances recommendation quality by harnessing redundant, rich, and similar domain data. Experiments validate the effectiveness of our method, demonstrating substantial improvements in recommendation quality.

1. INTRODUCTION

In the ever-expanding realm of online platforms and social networks, users engage across various channels, each presenting a distinct ecosystem. A critical challenge emerges from this fragmentation, where (*i*) multiple platforms offer the same products, (*ii*) users are active across various platforms, and (*iii*) data tends to be scarce, with some users or items receiving limited and insufficient ratings for effective learning. This leads to a fundamental question: Could we leverage the data across distinct yet conceptually similar platforms to enhance recommendation systems for all of them?

Cross-domain recommendation explores strategies to bridge the gap between different ecosystems. Some approaches adopt transfer learning-like strategies, while others focus on dual-target recommendations, addressing common users or items shared between domains. Some address the scenario where there are common users who interact in both domains, while others try to resolve the more general non-overlapping setting. Some make use of side information from user demographics or reviews, item metadata, and description; while others rely entirely on user-item interactions only. These categorical criteria split cross-domain recommendations into various settings (elaborated in Section 2).

In this paper, we address the specific challenge of cross-domain recommendation with *no* overlapping users, *no* overlapping items between the two platforms, and *no* side information available (NO3). Our proposed methodology establishes a bridge for learning across two related domains in recommendation systems, seeking to improve the accuracy and relevance of recommendations in scenarios where data sparsity poses a considerable hurdle. This approach is particularly crucial in instances where users interact on different platforms or social networks that share similarities, creating an opportunity to capitalize on the available data. As we delve into the details of our approach, it becomes evident that our methodology not only fills a critical gap in the existing literature but also lays the groundwork for more effective cross-domain recommendation systems in diverse online landscapes.

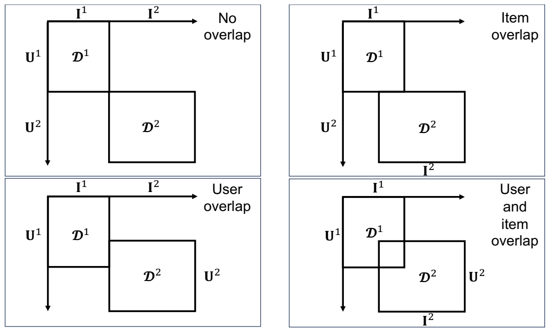

Figure 1: Four scenarios of overlapping user and item.

To this end, this paper makes several contributions. First, we introduce a novel problem setting tailored to address the challenge of scarce data in item-rich platforms, characterized by our unique setting (NO3-CDR) with no overlapping users, no overlapping items, and no side information. Second, to tackle the challenges, we propose two hard and soft user-matching learning algorithms, encapsulated in a bridge for learning across related domains in recommendation systems. Third, through empirical evidence from experiments, we demonstrate improvements in recommendation quality, offering a new perspective on cross-domain recommendation systems and alleviating privacy concerns by reducing the reliance on user identities. Our approach also leverages redundant data in similar domains to overcome scarce data hurdles in item-rich recommendation platforms.

2. PROBLEM FORMULATION

Cross-Domain Recommendation (CDR) in General. The original CDR is useful when data from one domain (known as the source domain), such as user-item interactions, is utilized to improve the recommendation process in a different but related domain (referred to as the target domain). The primary goal of CDR is to address challenges like data sparsity and the cold-start problem in the target domain by exploiting knowledge from the source domain.

Single-Target vs. Dual-Target Approaches. Prior research in CDR systems has explored methodologies aiming to transfer knowledge between distinct recommendation domains. Early work focuses on single-target approaches which typically entail exploiting redundant information from a source domain to a less abundant target domain. In these scenarios, the rich user or item information acquired from the source domain assists the learning process for the sparser target task. Techniques such as domain adaptation and transfer learning have been employed to improve recommendation performance specifically towards target domains.

Recently, there has been a growing interest in dual-target approaches, focusing on enhancing user and item recommendations across both domains. These methods seek to elevate recommendations by pinpointing and leveraging the common ground between user preferences and item attributes, thereby catering to the diverse interests of users across various domains.

Overlapping vs. Non-overlapping Data. Based on the overlap of users and items, cross-domain recommendations can be categorized into four scenarios as illustrated in Figure 1:

- *No overlap*: $\mathcal{U}^1 \cap \mathcal{U}^2 = \varnothing$ and $\mathcal{I}^1 \cap \mathcal{I}^2 = \varnothing$. There is no overlap between users and items.
- *User overlap*: $\mathcal{U}^1 \cap \mathcal{U}^2 \neq \varnothing$. There are shared users in both domains.
- *Item overlap*: $\mathcal{I}^1 \cap \mathcal{I}^2 \neq \varnothing$. There are shared items in both domains.
- *User and item overlap*: $\mathcal{U}^1 \cap \mathcal{U}^2 \neq \varnothing$ and $\mathcal{I}^1 \cap \mathcal{I}^2 \neq \varnothing$. There are overlaps between both users and items.

CDR with overlapped users/items seeks to capitalize on cross-domain information to enrich recommendations within the focal domain. Traditionally, such approaches presume users engaging across both domains, aiming to suggest source items to target users or mitigate cold-start issues for users new to the target domain. Yet, the constraint of overlapped users lacks practicality in the real world, considering that real user identities are not widely available.

Due to the limitations of assuming overlapped entities across domains, previous studies address the more general scenario of non-overlapping CDR, where they can leverage auxiliary information such as demographics and textual data across domains.

However, in scenarios where additional side information is unavailable or disregarded, the recommendation task relies solely on the historical user-item interactions. This situation poses challenges in bridging the gap between the two domains.

In this paper, we address the novel setting of *dual-target, non-overlapping, cross-domain recommendation, where auxiliary information is unavailable.* Our objective is to bridge the gap in user preferences between the two domains by aligning the underlying shared preferences of users across domains, distinguishing our novel problem setting from previous studies.

3. RELATED WORK

Cross-Domain Recommendation. CDR encompasses various problem settings.

Single-Target, Dual-Target, and Multi-Target. Foundational formulation of single-target setting [2; 3; 5] aims to mitigate data sparsity by utilizing redundant data or information from other domains to enhance the original domain. For instance, CBT [17] generates a codebook matrix to extract cluster-level ratings from an auxiliary domain to support the target domain. TALMUD [29] expands on this by incorporating multiple source domains with varying relevance rates. The research then extends to the multi-target CDR [6; 22; 30]. CLFM [6] adopts a multi-target approach, dividing the cluster-level codebook into common and domain-specific sections. RMGM [18] integrates multiple sparse domains sharing common latent cluster-level patterns into a generative model. Recently, dual-target CDR [35; 36] have gained more attention, aiming to improve recommendation quality across both domains. DTCDR [35] first formulates dual-target setting by sharing user knowledge across domains. GA-DTCDR [36] enhances this framework using graph and attention mechanisms to learn better representations of overlapping users.

User Overlapping. Full user overlap represents an extreme case where the same users exist across multiple domains [2; 12; 26], treating each domain as a vertical partition of the rating matrix. Techniques such as tri-factorization [12] and graph convolutional networks [7] are employed to align user preferences across domains. Conversely, [3; 6; 34] focus on the problem of *non-overlapping users*, leveraging user tags [3] and item features [31] as auxiliary information. Further research explores the concept of partial user overlap [28; 30; 37] using methods such as collective matrix factorization [30] and representation combination [35].

Using Side Information. Auxiliary knowledge, such as user tags [3; 34] and textual descriptions [15; 31], are also utilized to enhance recommendations.

Multi-Task Recommendation (MTRec). CDR can be viewed as a specific instance of MTRec, where similar or related tasks are learned concurrently across different domain datasets. Previous research in MTRec can be classified into three types: (*i*) parallel [8; 32], (*ii*) cascaded [33; 27], and (*iii*) auxiliary [10; 19].

In parallel MTRec, two or more recommendation tasks are optimized concurrently using a weighted sum of their losses. E.g., RnR [8] combines ranking and rating prediction tasks for personalized video recommendations, while MTER [32] integrates explanation generation alongside recommendation.

Cascaded multi-task recommendation refers to a sequential chain of tasks that must be performed in a strict order, modeling user behavior stages. An example in this domain is ESMM [27], which addresses sparsity and sample selection bias through an "impression $\rightarrow$ click $\rightarrow$ conversion" sequence.

In the auxiliary task relation, one task is designated as the main task, with other tasks serving as auxiliary tasks to enhance the main task's performance. This approach is similar to single-target cross-domain recommendation. MetaBalance [10] aims to reduce the gradient magnitude of auxiliary tasks to prioritize the target task objective, while MTRec [19] incorporates link prediction to support the primary recommendation task.

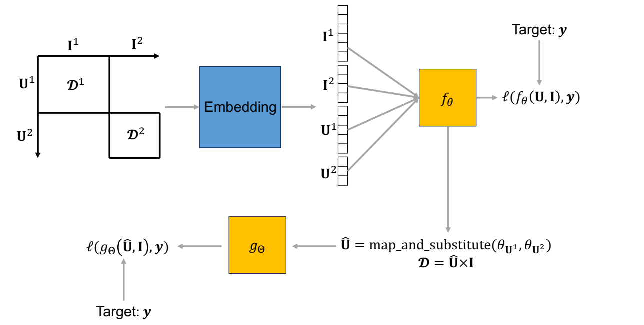

Figure 2: HNO3-CDR step-by-step workflow. Users, items, and ratings go through the embedding layer and recommendation model f_θ. Here, a generic recommender loss ℓ is computed by model prediction and target $\boldsymbol{y}$. Subsequently, based on the learned user representation, users from the two domains are mapped and substituted into new data $\mathcal{D}$. This new dataset is passed through a new recommendation model as an independent learning task.

Algorithm 1: HNO3-CDR Learning Algorithm

Input : $\mathcal{D}^1, \mathcal{D}^2, \mathcal{U} = \mathcal{U}^1 \cup \mathcal{U}^2, \mathcal{I} = \mathcal{I}^1 \cup \mathcal{I}^2$
$\theta^* = \arg\min_\theta \ell(\mathcal{D}_1 \cup \mathcal{D}_2 \mid \theta)$ ▷ Derive representations
$row\_ind, col\_ind = \text{Hungarian}(\theta^*_{\mathcal{U}^1}, \theta^*_{\mathcal{U}^2})$ ▷ Users matching
$\hat{\mathcal{U}} = \text{map}(\mathcal{U}, row\_ind, col\_ind)$ ▷ Mapping user indices
$\mathcal{D} \in \mathbb{R}^{|\hat{\mathcal{U}}| \times |\mathcal{I}|}$ ▷ New dataset from substituted users
$\Theta^* = \arg\min_\theta \ell(\mathcal{D} \mid \Theta)$ ▷ Learn until convergence
Output: Θ^*

4. METHODOLOGY

In the context of two distinct yet related tasks, $\mathcal{D}^1 \in \mathbb{R}^{|\mathcal{U}^1| \times |\mathcal{I}^1|}$ and $\mathcal{D}^2 \in \mathbb{R}^{|\mathcal{U}^2| \times |\mathcal{I}^2|}$, our objective is to develop a recommender model f parameterized by θ, denoted as f_θ, capable of capturing user preferences while enhancing recommendation performance for both tasks. Notably, we operate under the assumption that there is no predefined relationship between the sets of users ($\mathcal{U}^1$,$\mathcal{U}^2$) and items ($\mathcal{I}^1$,$\mathcal{I}^2$). Our focus is on the generalized scenario where user identities remain anonymous and cannot be directly mapped, and no additional item-related information, such as descriptions or reviews, is available.

Dual-target CDR. The dual-target framework is designed to optimize recommendation accuracy across domains. We aim to learn a unified model f_θ, that performs effectively in both domains:

$$\theta^* = \arg\min_\theta \Big(\ell(\mathcal{D}^1 \mid \theta) + \ell(\mathcal{D}^2 \mid \theta)\Big) \tag{1}$$

Here, ℓ represents a general model-agnostic loss function, such as Root Mean Squared Error (RMSE) for Matrix Factorization or Binary Cross-Entropy (BCE) for Neural Collaborative Filtering (NCF).

Optimizing vanilla dual-target CDR is equivalent to a simultaneous multi-task learning objective through a shared objective:

$$\theta^* = \arg\min_\theta \ell(\mathcal{D}^1, \mathcal{D}^2 \mid \theta) \tag{2}$$

In this scenario, the set of users, denoted as $\mathcal{U}$, is the union of two distinct individual user sets, i.e., $\mathcal{U} = \mathcal{U}^1 \cup \mathcal{U}^2$, with $|\mathcal{U}| = |\mathcal{U}^1| + |\mathcal{U}^2|$. Similarly, the set of items, denoted as $\mathcal{I}$, is the union of individual item sets, i.e., $\mathcal{I} = \mathcal{I}^1 \cup \mathcal{I}^2$, with $|\mathcal{I}| = |\mathcal{I}^1| + |\mathcal{I}^2|$.

4.1 HNO3-CDR: User Hard-Matching for Cross-Domain Recommendation

In the first attempt to bridge the connection of users in two domains, we find the hard-matching of every user from one domain to one corresponding user in the other domain, maximizing the similarities of matched users. Hungarian Algorithm [14] is a widely employed method to solve assignment problems. This classic algorithm minimizes the total cost of assignments in bipartite graphs, offering an efficient solution for various contexts. One user from the first domain can be assigned to at most one user in the other domain and vice versa. This results in a *hard* one-to-one user-matching across the two domains. Algorithm 1 and Figure 2 illustrate the step-by-step hard-matching learning algorithm for CDR. First, we obtain the optimal user representations from both domains in a multi-task learning setting, where the domain-specific datasets are combined as $\mathcal{D}^1 \cup \mathcal{D}^2$. The optimal parameters are learned by optimizing $\theta^* = \arg\min_\theta \ell(\mathcal{D}^1 \cup \mathcal{D}^2 \mid \theta)$. Next, we produce the mapping of the two user sets using the Hungarian algorithm. The resulting matching is then used to substitute users from one domain with their counterparts in the other. For example, if user $u_i^1 \in \mathcal{U}^1$ is matched with user $u_j^2 \in \mathcal{U}^2$, we replace u_j^2 with u_i^1. This creates a full overlapping user scenario, where the matched users are merged into a single unified set, denoted as $\hat{\mathcal{U}}$. Finally, using the substituted user set $\hat{\mathcal{U}}$, we construct a new dataset $\mathcal{D} \in \mathbb{R}^{|\hat{\mathcal{U}}| \times |\mathcal{I}|}$ and optimize a new model g_Θ accordingly.

4.2 SNO3-CDR: Soft-Matching End-To-End Cross-Domain Recommendation

HNO3-CDR faces several challenges. Firstly, it adopts a step-by-step learning process, where each step is executed discretely without a seamless flow, posing difficulties in optimization. Secondly, the mapping process occurs after the initial learning phase, creating uncertainty regarding the meaningfulness of the connection between the two user sets. Once this mapping is done, adjustments to enhance its suitability are not possible. To address these issues, we propose a solution that involves user soft-matching and functions as an end-to-end learning model. This model streamlines the learning process into a continuous flow and prioritizes

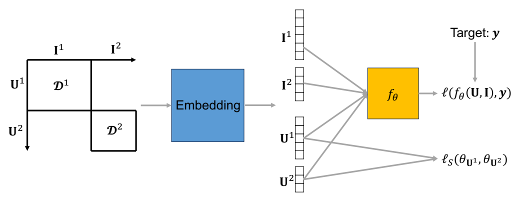

Figure 3: SNO3-CDR workflow. Users, items, and ratings go through the normal embedding layer and recommendation model f_θ to derive generic recommender loss ℓ between model prediction and target $\boldsymbol{y}$. Sinkhorn distance ℓ_S between two user sets acts as a bridge of users between the two domains and is combined with generic loss.

Algorithm 2: SNO3-CDR Learning Algorithm

Input : $\mathcal{D} = \mathcal{D}^1 \cup \mathcal{D}^2, \mathcal{U} = \mathcal{U}^1 \cup \mathcal{U}^2, \mathcal{I} = \mathcal{I}^1 \cup \mathcal{I}^2$, learning rate α

for *iteration = 1, . . . , w* **do**

$\quad \theta = \theta - \alpha \times \nabla \ell(\mathcal{D} \mid \theta)$ ▷ $w-$step warmup iterations

$\theta^* = \arg\min_\theta \ell(\mathcal{D} \mid \theta) + \gamma \ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2})$ ▷ Optimize until convergence

Output: θ^*

the optimization of general recommendations alongside the meaningful mapping of users. The end-to-end architecture ensures a continuous and adaptable mapping process, allowing for continuous enhancement of user representation with a focus on fostering meaningful connections throughout the model optimization process.

4.2.1 Sinkhorn distance

Optimal transport algorithms try to minimize transportation cost from *source/producer* to *target/consumer* given the producer' capacities and consumers' needs:

$$
\begin{aligned}
& d = \min \sum_{i,j} P_{i,j} C_{i,j} \\
\text{Subject to:} \quad & P_{i,j} \geq 0 \quad \text{for all } i, j \\
& \sum_j P_{i,j} = r_i \quad \text{for all } i \\
& \sum_i P_{i,j} = c_j \quad \text{for all } j
\end{aligned}
$$

where $P_{i,j}$ is the amount to transport from P_i to C_j, $C_{i,j}$ is cost to transport from P_i to C_j, r_i is capacity of P_i, and c_j is C_j's need.

Sinkhorn algorithm [1; 4] can be applied to transform the optimal transportation problem into the mapping of two "point clouds", where we transport "mass" from one set of points to another. [4] rewrites the original optimization formulation into Lagrange form:

$$
\begin{aligned}
d_S(P, C) = \sum_{i,j} P_{i,j} C_{i,j} - \frac{1}{\lambda} h(P) + \sum_i m_i \left(\sum_j P_{i,j} - r_i \right) \\
+ \sum_j n_j \left(\sum_i P_{i,j} - c_j \right)
\end{aligned} \tag{3}
$$

with m_i and n_j are Lagrange multipliers.

The derivative w.r.t. P can be easily derived by:

$$
\frac{\partial d_S}{\partial P_{i,j}} = C_{i,j} + \frac{1}{\lambda} + \frac{1}{\lambda} \log P_{i,j} + m_i + n_j
$$

This differentiable Sinkhorn distance can be seamlessly incorporated into any general objective of recommender models.

4.2.2 Mediate Latent Preferences by Sinkhorn Distance.

We constrain users from two domains to be close to each other without binding them tightly one-to-one. We define the Sinkhorn distance between two sets (i.e., point clouds) of user representations, $\mathcal{U}^1$ and $\mathcal{U}^2$, as:

$$
\ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2}) = d_S(\mathcal{U}^1, \mathcal{U}^2) + d_S(\mathcal{U}^2, \mathcal{U}^1) \tag{4}
$$

Here, $d_S(\mathcal{U}^1,\mathcal{U}^2)$ denotes the standard uni-directional Sinkhorn distance from point cloud $\mathcal{U}^1$ to $\mathcal{U}^2$, calculated using an arbitrary ground distance function (e.g., Euclidean, cosine) as the transportation cost between points in $\mathcal{U}^1$ and $\mathcal{U}^2$. This results in a symmetric, bi-directional distance measure. ℓ_S is differentiable with respect to both sets of representations, $\theta_{\mathcal{U}^1}$ and $\theta_{\mathcal{U}^2}$, making it suitable for gradient-based optimization within a recommender system framework. Alternatively, we could employ a standard uni-directional Sinkhorn distance, using either $d_S(\mathcal{U}^1,\mathcal{U}^2)$ or $d_S(\mathcal{U}^2,\mathcal{U}^1)$. Section 5 will show the impact of bi-directional and uni-directional formulations.

We incorporate ℓ_S into the training objective to mediate the latent preferences of users across domains. This encourages the user representations to be similar while retaining the capacity to capture domain-specific preferences. Conceptually, this can be formulated as a constrained optimization problem:

$$
\begin{aligned}
\theta^* &= \arg\min_\theta \ell(\mathcal{D}^1, \mathcal{D}^2 \mid \theta) \\
\text{Subject to:} \quad & \ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2}) \leq \alpha^2
\end{aligned}
$$

where $\ell(\mathcal{D}^1, \mathcal{D}^2 \mid \theta)$ is the primary recommendation loss function for data from domains $\mathcal{D}^1$ and $\mathcal{D}^2$, and α^2 is a positive tolerance threshold.

By rewriting the constraint as $\ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2}) - \alpha^2 \leq 0$, the final objective function for our end-to-end learning framework using the Lagrange multiplier is derived as:

$$
\begin{aligned}
\theta^* &= \arg\min_\theta \ell(\mathcal{D}^1, \mathcal{D}^2 \mid \theta) + \gamma \left(\ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2}) - \alpha^2 \right) \\
&\propto \arg\min_\theta \ell(\mathcal{D}^1, \mathcal{D}^2 \mid \theta) + \gamma \ell_S(\theta_{\mathcal{U}^1}, \theta_{\mathcal{U}^2})
\end{aligned} \tag{5}
$$

This augmented objective effectively balances the optimization of the primary recommendation task ℓ with the continuous and flexible mapping process ℓ_S, therefore promoting the transfer and adaptation of user preferences across domains by aligning their representations.

Alternatively, this augmented objective can be interpreted within a multi-task learning framework, where minimizing the transportation cost ℓ_s serves as an auxiliary task that supports and improves the performance of the primary recommendation task.

The overall learning process, illustrated in Algorithm 2 and Figure 3, involves an initial warm-up phase to learn meaningful user representations, followed by concurrently optimizing the augmented objective, which incorporates both the recommendation loss and the transportation cost.

5. EXPERIMENTS

Datasets. For experiments, the first three pairs of datasets are from Amazon[1]: *Books - Kindle Store*; *Electronics - Cell Phones and Accessories*; and *CDs and Vinyl - Digital Music*, chosen based on the assumption that users' preferences are likely shared between the two domains. For example, users who enjoy reading books may also be interested in similar Kindle e-books. To further diversify our analysis, we construct a fourth dataset from two sources: *Amazon Books - Book Crossing*[2], where the two share the same category of items but from different user sets and sources.

[1]https://nijianmo.github.io/amazon/index.html

[2]https://grouplens.org/datasets/book-crossing/

Table 1: Datasets stats for four scenarios

Dataset	Stats	Generic $\mathcal{D}^1$	Generic $\mathcal{D}^2$	Superset $\mathcal{D}^1$	Superset $\mathcal{D}^2$	Subset $\mathcal{D}^1$	Subset $\mathcal{D}^2$	Common $\mathcal{D}^1$	Common $\mathcal{D}^2$
Books	#ratings	8,898,041	982,619	8,898,041	967,196	1,319,803	982,619	1,319,803	967,196
—	#users	367,982	61,934	367,982	61,236	61,236	61,934	61,236	61,236
Kindle	#items	603,668	68,223	603,668	68,079	256,019	68,223	256,019	68,079
Electronics	#ratings	6,387,916	1,109,521	6,387,916	648,026	1,230,678	1,109,521	1,230,678	648,026
—	#users	694,953	154,813	694,953	81,381	81,381	154,813	81,381	81,381
Cell Phones	#items	157,693	47,607	157,693	46,996	134,621	47,607	134,621	46,996
CDs	#ratings	1,377,008	123,518	1,377,008	42,872	181,705	123,518	181,705	42,872
—	#users	107,546	12,381	107,546	3,720	3,720	12,381	3,720	3,720
Music	#items	71,943	9,906	71,943	9,113	49,898	9,906	49,898	9,113
AMZ Books	#ratings	223,302	197,140	-	-	-	-	-	-
—	#users	3,353	2,578	-	-	-	-	-	-
Book Crossing	#items	5,752	4,313	-	-	-	-	-	-

Four Scenarios. For comprehensive analysis, we explore four distinct scenarios, based on the overlap of two user sets $\mathcal{U}^1$ and $\mathcal{U}^2$, from the generic case with no constraint of users, to the extreme scenario where only users overlap between two domains are allowed, and the two middle ground scenarios.

- Scenario 1 (Generic): Any $\mathcal{U}^1$ and $\mathcal{U}^2$
- Scenario 2 (Superset): $\mathcal{U}^1 \supset \mathcal{U}^2$
- Scenario 3 (Subset): $\mathcal{U}^1 \subset \mathcal{U}^2$
- Scenario 4 (Common): $\mathcal{U}^1 = \mathcal{U}^2$

In all four cases, regardless of overlapping, *user identities are masked* so that the model treats the same user in two domains as two different users. Table 1 summarizes the respective statistics of the datasets under each of the four experimental scenarios.

Rating and Ranking Tasks. For evaluation, we employ two recommendation tasks: rating prediction and ranking prediction. We apply our model-agnostic proposed methods to two representative backbone models: Matrix Factorization (MF [13]) and Neural Collaborative Filtering (NCF [9]) and evaluate their performance. We use Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) for rating prediction, and Normalized Discounted Cumulative Gain (NDCG) and Recall with $k = 50$ for ranking prediction.

Comparative Methods. Due to its novel setting, there is no direct baseline for NO3-CDR. Previous dual-target cross-domain recommendation studies either (i) utilize shared parameters from the same users or items, which assumes user or item overlap–an assumption that does not hold in our setting–or (ii) leverage other data modalities as side information, which are also unavailable in our case. Therefore, we consider the comparative methods below:

- Base models: We use MF [13] and NCF [9] as backbone models for rating and ranking tasks, respectively. We combine data from two domains and train with one single model, with objective function in Equation 2.
- SinkhornCF [20]: Infuses Sinkhorn divergence of items into the learning objective. It can be applied to MF (i.e., SinkhornMF) and NCF (i.e., SinkhornNCF).
- NMF [16]: As recent studies [23; 24] suggest Non-negative Matrix Factorization (NMF) to be superior to the original MF, NMF is included as a baseline for rating prediction.
- VAECF [21] and its variants are widely used due to their non-linear probabilistic generative modeling. We include VAECF as a baseline for the ranking prediction task.

Table 2: The effects of aggregating user identities across domains for *Amazon CDs - Digital Music* dataset. Better results are in bold.

Training	CDs RMSE(↓)	CDs MAE(↓)	Digital Music RMSE(↓)	Digital Music MAE(↓)
Separately	0.6612	0.6103	0.5959	0.5729
Together	**0.6299**	**0.5775**	**0.5848**	**0.5621**

We adopt NMF and VAECF, which are considered superior to the backbone models MF and NCF, to evaluate whether the proposed methods can enhance the backbone models sufficiently to outperform these two baselines.

Hyper-parameter Tuning. Each dataset is partitioned into training, validation, and test sets using a chronological proportional split as described in prior works [11; 25], with a ratio of 60/20/20 for training, validation, and test sets. All methods are trained on the training set, tuned for optimal performance and model selection based on the validation set, and the best models are evaluated on the test set. We perform random search for hyper-parameter tuning, with the search space for some key hyper-parameters as follows: learning rate $\in [0.001, 0.1]$, embedding size $\in \{64, 100, 128, 256\}$, and control parameter $\gamma \in [0.1, 1.0]$. The number of warm-up iterations for SNO3 and HNO3-CDR is set to $w = 5$.

5.1 Research Questions (RQ) and Discussions

RQ1: The effects of using user identities across domains.

We first investigate the potential benefits of having user identities across domains. We carry out an experiment to compare the performance of training the model separately and together on the *CDs - Music* dataset. We first filter only users who have presented in both domains, then train the MF model in two different settings: (i) separately, where we train the model on each domain independently, and (ii) together, where we combine user-item interactions from both

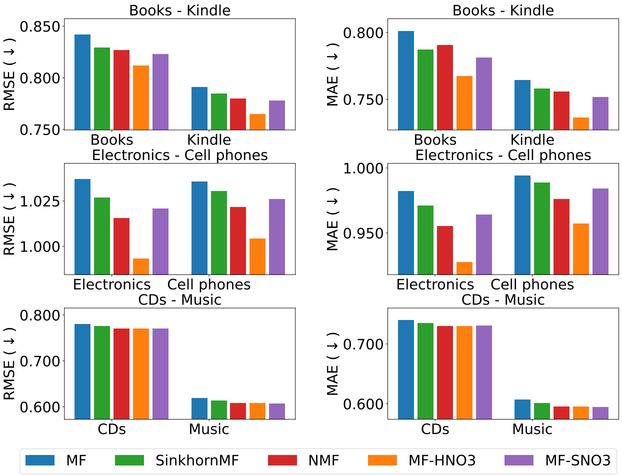

(a) Generic scenario.

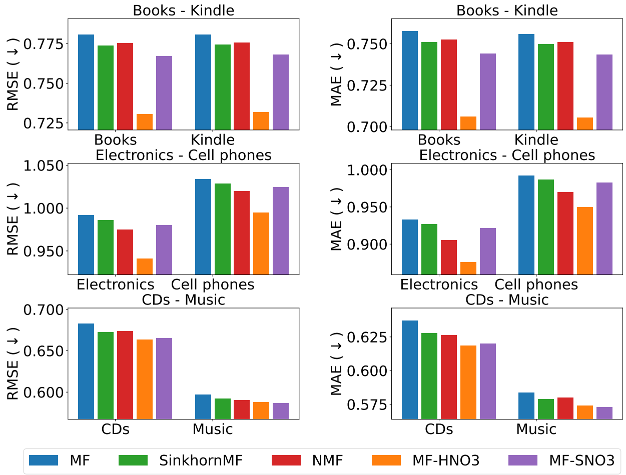

(b) Superset scenario.

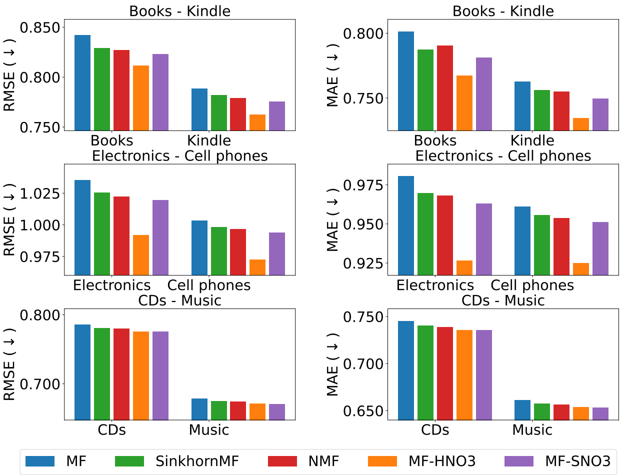

(c) Subset scenario.

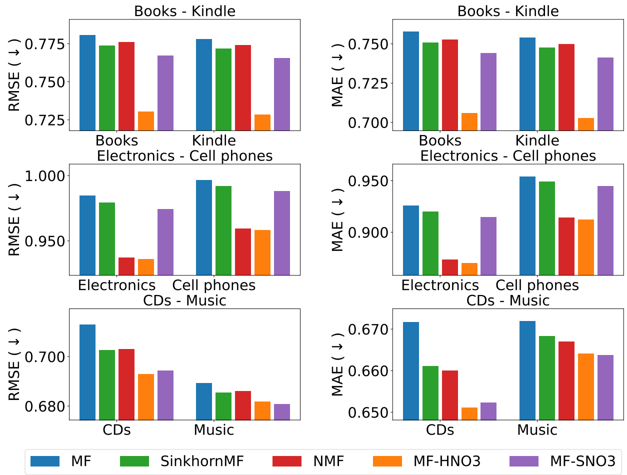

(d) Common scenario.

Figure 4: Rating prediction performances in four scenarios. For RMSE and MAE, the lower values (↓) indicate better results.

domains and train on the whole data.

Table 2 contrasts training the model separately versus jointly on the *CDs - Digital* dataset. The results show that joint training reduces both RMSE and MAE for *CDs* and *Music*, demonstrating improved performance over separate training. It aligns with the intuition that shared user identities can improve predictions across domains. Therefore, effective mechanisms for aligning user identities can be leveraged to enhance recommendations.

RQ2: How do the two variants NO3-CDR perform? Figures 4 and 5 present results across three *Amazon* datasets for two prediction tasks under four different scenarios. Comparing against benchmark baselines, we observe distinct behaviors in each task. For rating prediction (Figure 4), the MF-HNO3 and MF-SNO3 variants outperform SinkhornMF and NMF, both of which surpass traditional MF. Notably, MF-HNO3 consistently achieves the best performance, yielding significantly lower RMSE and MAE, followed by MF-SNO3 as the second-best performer. In contrast, for ranking prediction (Figure 5), the Hungarian-based NCF-HNO3 fails to surpass the NCF baseline, while SinkhornNCF and VAECF have superior performance over vanilla NCF. Among the NO3 variants, NCF-SNO3 consistently enhances the NCF backbone, achieving the best overall performance. It surpasses the two strongest NCF-based baselines in most cases, particularly in terms of NDCG and Recall. The only exceptions are NDCG on the *Books* domain (Figure 5a, top-left) and both NDCG and Recall on the *Electronics* and *Cell Phones* domains (Figure 5c, middle row), where SinkhornNCF marginally outperforms NCF-SNO3.

The choice between HNO3 and SNO3 depends on the specific recommendation task: HNO3 is more effective for rating prediction, while SNO3 excels in ranking tasks. This is likely due to the different ways the two backbone models generate item scores. In MF, ratings are directly predicted from user-item embeddings, which aligns well with the one-to-one matching of the HNO3 variant. In contrast, NCF generates user-item scores indirectly through multiple feedforward neural network layers, which benefits more from the flexible matching SNO3 for ranking tasks.

RQ3: The scenarios involving two different data sources.

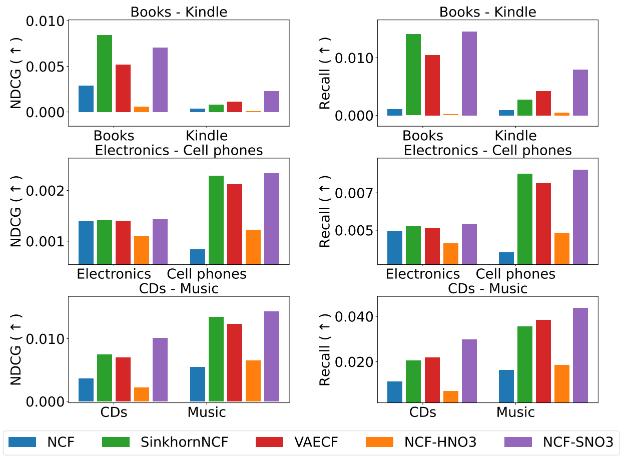

(a) Generic scenario.

(b) Superset scenario.

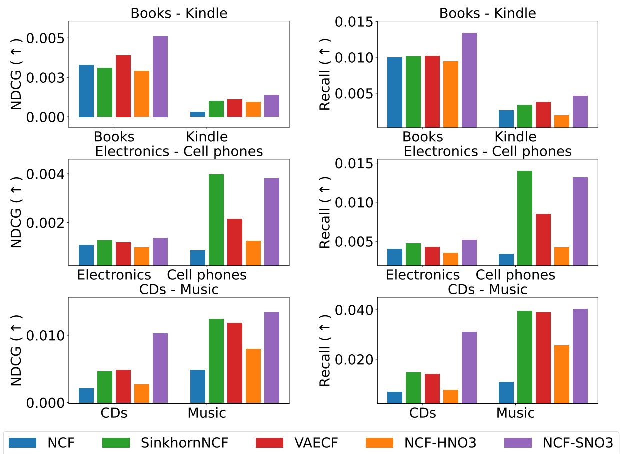

(c) Subset scenario.

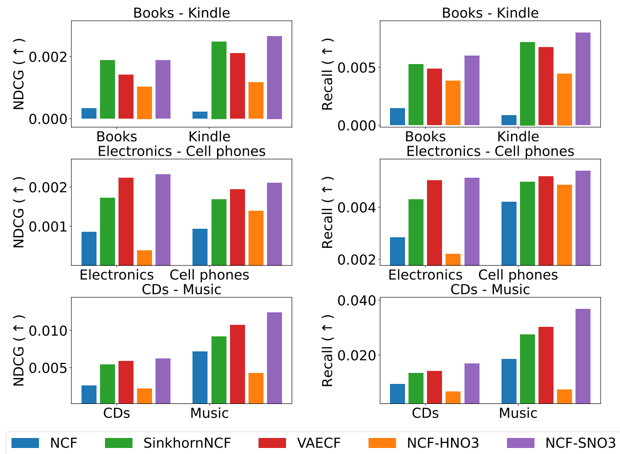

(d) Common scenario.

Figure 5: Ranking prediction performances in four scenarios. For NDCG and Recall, higher values (↑) indicate better results.

Table 4 presents the results of experiments conducted on datasets from two different sources: *Amazon Books* and *Book Crossing*. For the rating task, MF-HNO3 delivers the best performance in terms of RMSE and MAE, except for RMSE on *Amazon Books*, where it ranks second to NMF. SNO3-CDR closely follows behind. In the item ranking task, NCF-SNO3 outperforms the others in terms of NDCG and Recall, while NCF-HNO3 does not improve upon the NCF baseline. These findings align with previous results from the three *Amazon* dataset pairs. This supports the idea that aggregating data from multiple fragmented platforms can enhance performance. While more data does not always guarantee better results, effectively guiding the learning process allows the model to leverage richer information. The results also demonstrate the robustness of the proposed methods in improving recommendations across diverse data sources.

RQ4: Uni-directional versus bi-directional SNO3. SNO3-CDR offers the flexibility to transport bi-directionally between two "point clouds" $\mathcal{U}^1$ and $\mathcal{U}^2$. To see whether uni-directional or bi-directional yields superior recommendation, and whether there is an optimal assignment to each domain as source or target, we analyze three cases: (*i*) bi-directional transportation (i.e., no designated "target"), (*ii*) $\mathcal{U}^1$ as "target" point cloud, and (*iii*) $\mathcal{U}^2$ as "target" nodes.

Table 3 compares bi-directional and uni-directional MF-SNO3 and NCF-SNO3 across all four scenarios of the *CDs-Music* dataset. In all cases, the best uni-directional method outperforms the bi-directional method, improving results in both domains. No domain consistently outperforms the other. In three out of four scenarios, selecting one target domain enhances both rating and ranking predictions. The exception is the *Common* scenario: for ratings, selecting $\mathcal{D}_1$ as the target improves results, while for ranking, choosing $\mathcal{D}_2$ yields better performance.

In pursuit of optimal results for the one-sided SNO3, we propose an *automatic* method to identify the better "target" domain by selecting the domain with higher user representation variance. After warm-up epochs, we calculate and compare variances, choosing the domain with higher variance as the target. This *Auto* method achieves the best SNO3 results in most cases (see Table 3), except in the Superset scenario, where *Auto* performs better in $\mathcal{D}^1$ but not in $\mathcal{D}^2$. This discrepancy arises due to the extreme imbalance in dataset sizes (Table 1): $\mathcal{D}^1$ has over 1 million ratings, while $\mathcal{D}^2$ has only 42,872 ratings.

RQ5: Should we prioritize matching the same user across domains to enhance recommendation performance?

Table 3: Results of different “target” domain on *CDs - Music*’s four scenarios. Best results are in bold.

(a) Common scenario

Target domain	Rating prediction				Ranking prediction			
	D^1		D^2		D^1		D^2	
	RMSE (↓)	MAE (↓)	RMSE (↓)	MAE (↓)	NDCG (↑)	Recall (↑)	NDCG (↑)	Recall (↑)
None	0.7026	0.6611	0.6854	0.6683	0.0054	0.0134	0.0092	0.0274
$\mathcal{D}^1$	**0.6942**	**0.6523**	**0.6808**	**0.6637**	0.0065	0.0164	0.0082	0.0237
$\mathcal{D}^2$	0.7129	0.6716	0.6892	0.6719	**0.0069**	**0.0169**	**0.0121**	**0.0360**
Auto	**0.6942**	**0.6523**	**0.6808**	**0.6637**	**0.0069**	**0.0169**	**0.0121**	**0.0360**

(b) Superset scenario

Target domain	Rating prediction				Ranking prediction			
	D^1		D^2		D^1		D^2	
	RMSE (↓)	MAE (↓)	RMSE (↓)	MAE (↓)	NDCG (↑)	Recall (↑)	NDCG (↑)	Recall (↑)
None	0.6103	0.5900	0.5922	0.5790	0.0045	0.0134	0.0124	0.0386
$\mathcal{D}^1$	**0.6041**	**0.5836**	**0.5864**	**0.5731**	**0.0046**	**0.0146**	0.0135	0.0395
$\mathcal{D}^2$	0.6161	0.5958	0.5969	0.5837	0.0028	0.0085	0.0135	**0.0414**
Auto	**0.6041**	**0.5836**	**0.5864**	**0.5731**	**0.0046**	**0.0146**	0.0135	0.0395

(c) Subset scenario

Target domain	Rating prediction				Ranking prediction			
	D^1		D^2		D^1		D^2	
	RMSE (↓)	MAE (↓)	RMSE (↓)	MAE (↓)	NDCG (↑)	Recall (↑)	NDCG (↑)	Recall (↑)
None	0.6726	0.6277	0.5922	0.5790	0.0035	0.0082	0.0082	0.0234
$\mathcal{D}^1$	**0.6650**	**0.6200**	**0.5864**	**0.5731**	**0.0036**	**0.0089**	**0.0095**	**0.0237**
$\mathcal{D}^2$	0.6824	0.6373	0.5969	0.5837	0.0023	0.0071	0.0048	0.0217
Auto	**0.6650**	**0.6200**	**0.5864**	**0.5731**	**0.0036**	**0.0089**	**0.0095**	**0.0237**

(d) Generic scenario

Target domain	Rating prediction				Ranking prediction			
	D^1		D^2		D^1		D^2	
	RMSE (↓)	MAE (↓)	RMSE (↓)	MAE (↓)	NDCG (↑)	Recall (↑)	NDCG (↑)	Recall (↑)
None	0.7757	0.7349	0.6138	0.6014	0.0035	0.0107	0.0047	0.0130
$\mathcal{D}^1$	0.7806	0.7394	0.6192	0.6068	0.0025	0.0085	0.0044	0.0110
$\mathcal{D}^2$	**0.7709**	**0.7301**	**0.6074**	**0.5948**	**0.0036**	**0.0112**	**0.0055**	**0.0163**
Auto	**0.7709**	**0.7301**	**0.6074**	**0.5948**	**0.0036**	**0.0112**	**0.0055**	**0.0163**

Table 4: Results for *Amazon Books - Book Crossing* dataset. Note that in Amazon Books, the rating scale is from 1 to 5, while for Book Crossing is from 1 to 10. Best results are in bold, while second-best results are in *italic*.

(a) Rating Prediction

Model	AMZ Books		Book Crossing	
	RMSE	MAE	RMSE	MAE
MF	0.9080	0.8429	3.3754	3.1218
SinkhornMF	0.8878	0.8289	3.3116	3.1498
NMF	**0.8853**	0.8226	3.2920	3.0703
MF-HNO3	*0.8864*	**0.8193**	**3.2564**	**3.0700**
MF-SNO3	0.8865	*0.8195*	*3.2771*	*3.0701*

(b) Ranking Prediction

Model	AMZ Books		Book Crossing	
	NDCG (%)	Recall (%)	NDCG (%)	Recall (%)
NCF	*0.1075*	*0.3494*	0.1324	0.3268
SinkhornNCF	0.0938	0.2784	*0.1667*	*0.3582*
VAECF	0.0786	0.2310	0.1308	0.3453
NCF-HNO3	0.0890	0.2069	0.1168	0.2954
NCF-SNO3	**0.1198**	**0.3709**	**0.1709**	**0.3838**

Table 5: Case study in CDs-Music dataset

User in CDs: A117WAVHO1WAIE		User in Music: A8QZWK9SUH66P	
Items rated	Items categories	Items rated	Items categories
The Commodores	R&B, Funk, Pop	Doo-Wops & Hooligans	Pop, R&B
Earth Wind & Fire	R&B, Funk, Soul	Waking Up	Pop, Rock
Song of Solomon	Rock, Pop	X	Pop, R&B
Carpenters Gold	Pop	Here’s To The Good Times	Pop, Rock
Piano Prophet	Jazz, R&B	The Fault In Our Stars	Rock
		The Hunting Party	Rock

User in CDs: A28DBLK5JB17P3		User in Music: A167KI3P7XN1AM	
Items rated	Items categories	Items rated	Items categories
Led Zeppelin: Box	Rock, Metal		
Led Zeppelin I	Rock, Metal	Made In The A.M.	Pop, Rock
Led Zeppelin II	Rock, Metal		
Houses of the Holy	Rock, Metal	Somewhere In Time LP	Rock, Metal
At Your Service	Pop, Rock		

User in CDs: A28DBLK5JB17P3		User in Music: A1VFOUHOYX29YP	
Items rated	Items categories	Items rated	Items categories
Led Zeppelin: Box	Rock, Metal	Light Me Up	Rock, Metal
Led Zeppelin I	Rock, Metal	Hit Me Like A Man	Rock, Metal
Led Zeppelin II	Rock, Metal	Bad Magic - Motörhead	Rock, Metal
Houses of the Holy	Rock, Metal	Dystopia - Megadeth	Rock, Metal
At Your Service	Pop, Rock	XI Metal - Church	Rock, Metal

Users may portray different preferences across platforms, such as purchasing classical music on *CDs and Vinyl* and modern trending songs on *Digital Music*. Our goal is to enhance recommendations on both platforms rather than focusing solely on matching users across domains, as we assume no overlap in users.

However, *though not used in the learning as presumed non-existent*, the availability of user identity information allows us to investigate whether the algorithms match users across domains correctly. We investigate the user mapping accuracy in *CDs - Music* dataset’s Common scenario, using MF-HNO3, since it performs best in rating prediction; and NCF-SNO3 for ranking. Surprisingly, out of 3,720 users across both domains, MF-HNO3 accurately maps only 1 to 3 users on different runs. While NCF-SNO3 does not output user mapping, we derive the mapping based on the closest Sinkhorn distances of final user representations, and the result is 0 to 3 correct user pairs.

HNO3 is a step-by-step learning process and mapping quality solely relies on user representation derived from the initial learning model. For SNO3, the control variable γ in Equation 5 can be adjusted to balance recommendation and transportation objectives. However, as γ increases (favoring user mapping), recommendation performance gradually decreases. The Sinkhorn distance in SNO3 acts as a flexible bridge between domains, where matching users is not prioritized to achieve the best recommendation quality.

5.2 Case Study: Example Matched User Pairs

Table 5 presents three user pairs from the *CDs* domain along with their corresponding matches from the *Music* domain.

In the first pair, both users show similar preferences for a mix of R&B, Pop, and Rock. User A117WAVHO1WAIE has a diverse taste, enjoying artists like *The Commodores*, *Earth Wind & Fire*, and *The Carpenters*, ranging from classic R&B and funk to pop. Interestingly, her match in the music domain, user A8QZWK9SUH66P, also appreciates Pop and R&B, with selections like *Bruno Mars' "Doo-Wops & Hooligans"* and *Florida Georgia Line's "Here's To The Good Times"*, showcasing a similar inclination to pop and rock.

The second pair, user A28DBLK5JB17P3 in *CDs* and user A167KI3P7XN1AM in *Music*, exhibited more distinct common preferences. They are deeply rooted in rock and metal, especially classic metal rock. In the third pair, user from the second pair, A28DBLK5JB17P3, is also the best match for the user in *Music*, A1VFOUHOYX29YP, who also roots for rock albums, such as The Pretty Reckless' *"Light Me Up"* and *"Hit Me Like A Man"*.

NCF-SNO3 effectively captures the similarities among intricate user preferences. The consolidation of these identified parallels among matched user pairs serves to reinforce the notion of preference bridging, rather than prioritizing the enhancement of correct matching accuracy. While the optimal match for a user across domains may not fully align with their unique preferences, they may exhibit a greater degree of similarity in their preferences compared to their own preferences in different domains.

6. CONCLUSION

This paper addresses the challenge of scarce data in recommendation systems. We introduce the novel scenario of NO3-CDR framework and propose a unique approach to enhance recommendation systems by leveraging connections across distinct yet conceptually similar datasets from multiple platforms based on user underlying preferences. Our methodology focuses on bridging the gap between these platforms, enabling mutual improvements in recommendation quality while respecting user privacy. Empirical experiments demonstrate the effectiveness of our approach in improving recommendation quality, showcasing its potential to address data scarcity challenges in fragmented cross-domain recommendation systems.

7. REFERENCES

[1] J. Altschuler, J. Niles-Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. *Advances in neural information processing systems*, 30, 2017.

[2] S. Berkovsky, T. Kuflik, and F. Ricci. Cross-domain mediation in collaborative filtering. In *User Modeling 2007: 11th International Conference, UM 2007, Corfu, Greece, July 25-29, 2007. Proceedings 11*, pages 355–359. Springer, 2007.

[3] W. Chen, W. Hsu, and M. L. Lee. Making recommendations from multiple domains. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 892–900, 2013.

[4] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

[5] A. M. Elkahky, Y. Song, and X. He. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *Proceedings of the 24th international conference on world wide web*, pages 278–288, 2015.

[6] S. Gao, H. Luo, D. Chen, S. Li, P. Gallinari, and J. Guo. Cross-domain recommendation via cluster-level latent factor model. In *Proceedings of the 2013th European Conference on Machine Learning and Knowledge Discovery in Databases-Volume Part II*, pages 161–176, 2013.

[7] L. Guo, L. Tang, T. Chen, L. Zhu, Q. V. H. Nguyen, and H. Yin. Da-gcn: A domain-aware attentive graph convolution network for shared-account cross-domain sequential recommendation. *arXiv preprint arXiv:2105.03300*, 2021.

[8] G. Hadash, O. S. Shalom, and R. Osadchy. Rank and rate: multi-task learning for recommender systems. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pages 451–454, 2018.

[9] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pages 173–182, 2017.

[10] Y. He, X. Feng, C. Cheng, G. Ji, Y. Guo, and J. Caverlee. Metabalance: improving multi-task recommendations via adapting gradient magnitudes of auxiliary tasks. In *Proceedings of the ACM Web Conference 2022*, pages 2205–2215, 2022.

[11] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl. An algorithmic framework for performing collaborative filtering. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 230–237, 1999.

[12] L. Hu, J. Cao, G. Xu, L. Cao, Z. Gu, and C. Zhu. Personalized recommendation via cross-domain triadic factorization. In *Proceedings of the 22nd international conference on World Wide Web*, pages 595–606, 2013.

[13] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.

[14] H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[15] A. Kumar, N. Kumar, M. Hussain, S. Chaudhury, and S. Agarwal. Semantic clustering-based cross-domain recommendation. In *2014 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pages 137–141. IEEE, 2014.

[16] D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. *Advances in neural information processing systems*, 13, 2000.

[17] B. Li, Q. Yang, and X. Xue. Can movies and books collaborate? cross-domain collaborative filtering for sparsity reduction. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 2052–2057, 2009.

[18] B. Li, Q. Yang, and X. Xue. Transfer learning for collaborative filtering via a rating-matrix generative model. In *Proceedings of the 26th annual international conference on machine learning*, pages 617–624, 2009.

[19] H. Li, Y. Wang, Z. Lyu, and J. Shi. Multi-task learning for recommendation over heterogeneous information network. *IEEE Transactions on Knowledge and Data Engineering*, 34(2):789–802, 2020.

[20] X. Li, J. Y. Chin, Y. Chen, and G. Cong. Sinkhorn collaborative filtering. In *Proceedings of the web conference 2021*, pages 582–592, 2021.

[21] D. Liang, R. G. Krishnan, M. D. Hoffman, and T. Jebara. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*, pages 689–698, 2018.

[22] M. Liu, J. Li, G. Li, and P. Pan. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pages 885–894, 2020.

[23] X. Luo, M. Zhou, Y. Xia, and Q. Zhu. An efficient non-negative matrix-factorization-based approach to collaborative filtering for recommender systems. *IEEE Transactions on Industrial Informatics*, 10(2):1273–1284, 2014.

[24] X. Luo, Y. Zhou, Z. Liu, L. Hu, and M. Zhou. Generalized nesterov's acceleration-incorporated, non-negative and adaptive latent factor analysis. *IEEE Transactions on Services Computing*, 15(5):2809–2823, 2021.

[25] H. Ma, I. King, and M. R. Lyu. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 203–210, 2009.

[26] M. Ma, P. Ren, Y. Lin, Z. Chen, J. Ma, and M. d. Rijke. 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*, pages 685–694, 2019.

[27] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, and K. Gai. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1137–1140, 2018.

[28] T. Man, H. Shen, X. Jin, and X. Cheng. Cross-domain recommendation: An embedding and mapping approach. In *IJCAI*, volume 17, pages 2464–2470, 2017.

[29] O. Moreno, B. Shapira, L. Rokach, and G. Shani. Talmud: transfer learning for multiple domains. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 425–434, 2012.

[30] D. Rafailidis and F. Crestani. Top-n recommendation via joint cross-domain user clustering and similarity learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 426–441. Springer, 2016.

[31] K. Shu, S. Wang, J. Tang, Y. Wang, and H. Liu. Crossfire: Cross media joint friend and item recommendations. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 522–530, 2018.

[32] N. Wang, H. Wang, Y. Jia, and Y. Yin. Explainable recommendation via multi-task learning in opinionated text data. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 165–174, 2018.

[33] D. Xi, Z. Chen, P. Yan, Y. Zhang, Y. Zhu, F. Zhuang, and Y. Chen. Modeling the sequential dependence among audience multi-step conversions with multi-task learning in targeted display advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3745–3755, 2021.

[34] D. Yang, J. He, H. Qin, Y. Xiao, and W. Wang. A graph-based recommendation across heterogeneous domains. In *proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 463–472, 2015.

[35] F. Zhu, C. Chen, Y. Wang, G. Liu, and X. Zheng. Dtcdr: A framework for dual-target cross-domain recommendation. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1533–1542, 2019.

[36] F. Zhu, Y. Wang, C. Chen, G. Liu, and X. Zheng. A graphical and attentional framework for dual-target cross-domain recommendation. In *IJCAI*, volume 21, page 39, 2020.

[37] Y. Zhu, Z. Tang, Y. Liu, F. Zhuang, R. Xie, X. Zhang, L. Lin, and Q. He. Personalized transfer of user preferences for cross-domain recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pages 1507–1515, 2022.