

MIT Open Access Articles

Cross-Domain Latent Factors Sharing via Implicit Matrix Factorization

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Samra, Abdulaziz, Frolov, Evgeny, Vasilev, Alexey, Grigorevskiy, Alexander and Vakhrushev, Anton. 2024. "Cross-Domain Latent Factors Sharing via Implicit Matrix Factorization."

As Published: <https://doi.org/10.1145/3640457.3688143>

Publisher: ACM|18th ACM Conference on Recommender Systems

Persistent URL: <https://hdl.handle.net/1721.1/159402>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Cross-Domain Latent Factors Sharing via Implicit Matrix Factorization

Abdulaziz Samra
abdulaziz.samra@skoltech.ru
Skolkovo Institute of Science and
Technology
Moscow, Russia

Evgeney Frolov
frolov@airi.net
AIRI, Skolkovo Institute of Science
and Technology
Moscow, Russia

Alexey Vasilev
alexvl.vasilev@yandex.ru
Sber AI Lab
Moscow, Russia

Alexander Grigorevskiy
alex.grigorievskiy@gmail.com
Independent Researcher
Helsinki, Finland

Anton Vakhrushev
btbpanda@gmail.com
Independent Researcher
Moscow, Russia

ABSTRACT

Data sparsity has been one of the long-standing problems for recommender systems. One of the solutions to mitigate this issue is to exploit knowledge available in other source domains. However, many cross-domain recommender systems introduce a complex architecture that makes them less scalable in practice. On the other hand, matrix factorization methods are still considered to be strong baselines for single-domain recommendations. In this paper, we introduce the CDIMF, a model that extends the standard implicit matrix factorization with ALS to cross-domain scenarios. We apply the Alternating Direction Method of Multipliers to learn shared latent factors for overlapped users while factorizing the interaction matrix. In a dual-domain setting, experiments on industrial datasets demonstrate a competing performance of CDIMF for both cold-start and warm-start. The proposed model can outperform most other recent cross-domain and single-domain models. We also provide the code to reproduce experiments on GitHub.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Collaborative filtering**.

KEYWORDS

Alternating Direction Method of Multipliers (ADMM), Alternating Least Squares (ALS), Cross-Domain Recommender System (CDRS), Implicit Matrix Factorization

ACM Reference Format:

Abdulaziz Samra, Evgeney Frolov, Alexey Vasilev, Alexander Grigorevskiy, and Anton Vakhrushev. 2024. Cross-Domain Latent Factors Sharing via Implicit Matrix Factorization. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3640457.3688143>

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

RecSys '24, October 14–18, 2024, Bari, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0505-2/24/10
<https://doi.org/10.1145/3640457.3688143>

1 INTRODUCTION

In recent decades, recommender systems (RSs) have become increasingly popular in many online markets and social media platforms. They generate personalized suggestions for new content for users, providing them with a more engaging experience. Among recommendation algorithms, both Matrix Factorization (MF) and Neural Network approaches remain applicable in the industry. Recent studies have shown that MF algorithms are fairly strong baselines, outperforming many neural network algorithms [8]. The use of MF algorithms is determined by their ease of use and speed of inference. However, such algorithms are not free of drawbacks. Data sparsity [18] and the cold start [41] are the common sources of problems that easily impede learning proper user and item representations and severely influence recommendation accuracy.

Large tech companies such as Amazon, Google, Facebook, and many others are actively developing their ecosystems which consist of groups of interconnected companies [16]. These companies complement each other trying to cover users' interests in all item niches. To compete in the market, ecosystems include new organizations covering niches that have not yet been occupied. This renders the cold-start issue, i.e., the complete lack of user-item interactions, or sparsity issue due to a limited number of users with a narrow interactions history. These issues combined form one of the most challenging problems facing RSs in different companies, related to insufficient transactional and feedback data required for accurate inference of users' similarities [14].

Cross-domain recommendation (CDR) was proposed some years ago [40] to retain and expand the user base. The recommendation systems organizations can use data from other domains, such as related companies from the same ecosystem with a more extended history. Cross-domain recommender systems (CDRS) aim to increase the quality of recommendations by exploiting knowledge from multiple source domains [2]. This type of recommender system can help either when one domain is rich in information and the other has more sparse data (single target CDR) or when they are equally sparse (dual target CDR) [40]. In this setting, it is assumed that there is a shared set of users between related domains. However, some privacy restrictions exist for sharing user data across multiple domains in such scenarios.

Privacy concerns set a new challenge for cross-domain recommender systems. Cross-domain information transfer must occur

without explicitly exposing users' history to the other domains. The shared information must be safe to transmit over insecure channels so that user data can not be extracted. We also considered this requirement in our method, and propose to exchange a mixture of two different matrices between domains, making restoring the original data difficult.

On the other hand, a significant drawback of modern CDR methods is the requirement for resources and the complexity of implementation. Thus, we propose CDIMF; a cross-domain variant of the alternating least-squares (ALS) [13] which is well-known for its simple implementation and fast convergence.

Our main contributions to this work can be summarized as follows:

- We proposed the Cross-Domain Implicit Matrix Factorization (CDIMF) model for cross-domain recommendations based on the ALS model.
- The proposed algorithm outperformed state-of-the-art approaches on different pairs of datasets.
- The new approach can be used for the warm-start scenario, as well as for the cold-start scenario.

The structure of the paper is organized as follows. In Section 2, we articulate the problem statement we aim to address. Section 3 contextualizes our work within the existing body of research on CDR. Our proposed model is introduced in Section 4 while the experimental framework, including methodology and setup, is elucidated in Section 5. Section 6 presents the empirical results, including comparisons with existing models and an analysis of our model's performance. Finally, Section 7 provides a comprehensive discussion of our findings, conclusions drawn, and prospects for future research.

2 PROBLEM STATEMENT

One of the famous models to predict user feedback is to learn latent factors for users and items. Let's consider the basic matrix factorization problem:

$$\min_{\{X\},\{Y\}} F(X, Y), \quad (1)$$

where:

$$F(X, Y) = \|C \circ (P - XY^T)\|_F^2 + \lambda \left(\|\Lambda_X X\|_F^2 + \|\Lambda_Y Y\|_F^2 \right). \quad (2)$$

In this notation, $P \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$ is the user-item interaction matrix, $X \in \mathbb{R}^{|\mathcal{U}| \times d}$ and $Y \in \mathbb{R}^{|\mathcal{I}| \times d}$ are the latent factors of size d for users from the set \mathcal{U} and items from the set \mathcal{I} , respectively. The weighting matrix C is usually used to add different levels of importance to the errors $(P - XY^T)$ in the loss function. We add the norms of the latent factors to the objective function for regularization. The hyperparameter λ controls the regularization effect globally, and Λ_X, Λ_Y are diagonal matrices used to regularize each user/item embedding individually. More details about our selection for values of Λ_X, Λ_Y , and C are in Section 4.3.

We aim to extend the formula (2) to the case of N domains, the minimization problem would be:

$$\min_{\{X_i\},\{Y_i\}} \sum_i^N F_i(X_i, Y_i), \quad (3)$$

where the subscript i refers to the corresponding matrices on each domain. In cross-domain scenarios, having a set of overlapped users between domains can unveil hidden correlations that can help to learn a better recommendation model. For this purpose, we introduce equality constraints for all user factors across domains

$$\begin{aligned} \min_{\{X_i\},\{Y_i\}} \sum_i^N F_i(X_i, Y_i), \\ \text{s.t. } X_1 = X_2 = \dots = X_N. \end{aligned} \quad (4)$$

The parameters of this objective function $\{X_i\}, \{Y_i\}$ are distributed over the N domains, and that function can be jointly minimized. We hypothesize that the equality constraints would help with information sharing between domains by making the factorization process on each domain aware of the pattern in the others.

3 RELATED WORK

Implicit Matrix factorization (IMF) model with alternating least-squares (ALS) algorithm [13] has been a widely popular recommendation model in academic research [6], [11]. It has also been a workhorse model in many industrial applications due to its simplicity, scalability, and strong performance [7]. Despite the invention of more powerful deep-learning models, IMF has proven to be a strong baseline model for collaborative filtering recommender systems, provided hyper-parameters are properly tuned [27]. However, it implies that all the user data is located on one server (or cluster) and that data is coming from a single domain. Our algorithm addresses these ALS issues. It allows each data provider to keep the user data private and exchange only perturbed user embedding representations. Moreover, by putting different weights on different data domains, the model extends to the CDR setup.

The Alternating Direction Method of Multipliers (ADMM) has been successfully applied in recommender systems design before. For instance, [31] considers ADMM as a convenient optimization algorithm for the SLIM [22] model which is a sparse linear autoencoder for collaborative filtering. Research in [36], [34], [38] apply ADMM to optimize IMF objective function in a high-performance computing environment. Hence, their goal is to address the large-scale problem and to shorten the model training time. Furthermore, [34], [38] utilize the differential privacy technique to prevent uncovering user identities. However, the model is still trained on a centralized cluster, which implies access to the entire dataset. Overall, previous usage of ADMM is unrelated to cross-domain setup and restrictions on sharing user interaction data.

A recent overview of CDRS is provided in [37]. The proposed taxonomy of methods consists of two levels. The first level comprises various degrees of user data overlap between domains. The second level considers different recommendation tasks that can be posed on CDRS. On the first taxonomy level, our formulation is applicable when there is no item overlap and full user overlap between domains. Following the second level, our formulation can be considered as both intra-domain and inter-domain formulation. Also, it is a multi-target system since recommendations can be made for all involved domains.

Before the deep learning era, CDRS was mostly based on linear matrix factorization models. For example, work [23] proposes an IMF model to learn user and item embeddings on the source

domain. Target domain embeddings are found to be close to the source embeddings in a Frobenious norm sense. In the article [21] a linear (also nonlinear) mapping function between source and target domain embeddings is proposed. In contrast, our algorithm does not separate source and target domains thereby being a multi-target algorithm. Collective matrix factorization (CMF) based algorithms [30], [39], [37] has been also applied to the CDR problem. In contrast to our approach, Stochastic Gradient Descent is typically used for solving the CMF optimization problem. Our algorithm is an extension of the ALS algorithm which makes it familiar to the broad academia and industry communities. It is capable of reproducing the exact single-domain ALS solution when hyperparameters are set properly.

A vast amount of deep learning models have been proposed for CDR [37]. Motivated by the recent work of Cao et al. [3], who proposed a universal CDR model, we aim to develop a model that can be applied to multiple CDR scenarios. Unlike UniCDR, which adopts a neural contrastive learning approach, we will extend a standard MF model to the cross-domain setting to inherit the advantages of simplicity and computational efficiency.

4 PROPOSED MODEL

To achieve knowledge transfer across domains, we aim to make the factors of the shared users equal over all domains at the end of training. A famous method to train a distributed model with shared parameters is applying the Alternating Direction Method of Multipliers (ADMM). It is an algorithm that is used to solve convex optimization problems. It works by breaking the problem into smaller subproblems that can be solved on separate sides [1]. ADMM sharing optimization problem is:

$$\begin{aligned} \min_{x_i, z} \sum_{i=1}^N f_i(x_i) + g(z), \\ \text{s.t. } x_i - z = 0. \end{aligned} \quad (5)$$

This problem is called *Global Variable Consensus with Regularization* [1]. It is about minimizing the functions $\{f_i(x)\}$ located on different sides while trying to keep the values of $\{x_i\}$ as close as possible to the shared variable z . Iterations for solving this problem are defined as follows:

$$x_i^{k+1} = \arg \min_x f_i(x) + \frac{\rho}{2} \|x - z^k + u_i^k\|_2^2, \quad (6)$$

$$z^{k+1} = \arg \min_z g(z) + \frac{N\rho}{2} \|z - \langle x \rangle^{k+1} - \langle u \rangle^k\|_2^2, \quad (7)$$

$$u_i^{k+1} = u_i^k + (x_i^{k+1} - z^{k+1}). \quad (8)$$

In these equations, k is the iteration number, u_i is the dual variable related to the equality constraint $x_i - z = 0$, and ρ is the penalty parameter that controls how strong the local variable x_i is attracted to the global z in every iteration.

Our proposed approach is to get common user factors over domains while factorizing interaction matrices. Then, we rewrite the optimization problem (10) as:

$$\begin{aligned} \min_{\{X_i\}, \{Y_i\}, Z} \sum_i^N F_i(X_i, Y_i) + g(Z), \\ \text{s.t. } X_i - Z = 0. \end{aligned} \quad (9)$$

4.1 Update Iterations for CDIMF

Let us write the Augmented Lagrangian for the optimization problem (9)

$$\mathcal{L}(\{X_i\}, \{Y_i\}, Z, \{U_i\}) = F_i(X_i, Y_i) + \frac{\rho}{2} \|X_i - Z + U_i\|_F^2 + g(Z). \quad (10)$$

As mentioned before, ρ is a penalty parameter. Experiments (Section 6.2.1) will demonstrate that this parameter controls the speed of information sharing. Similarly, variables $\{U_i\}$ are the dual variables for the equality constraints $X_i - Z = 0$. The iterations for solving the federated matrix factorization problem using ADMM, are:

$$X_i^{k+1}, Y_i^{k+1} = \arg \min_{X_i, Y_i} F_i(X_i, Y_i) + \frac{\rho}{2} \|X_i - Z^k + U_i^k\|_F^2, \quad (11)$$

$$Z^{k+1} = \text{prox}_{g, \frac{1}{\rho N}} \left(\langle X \rangle^{k+1} + \langle U \rangle^k \right), \quad (12)$$

$$U_i^{k+1} = U_i^k + (X_i^{k+1} - Z^{k+1}). \quad (13)$$

Here, N is the total number of domains, and $\text{prox}_{g, \mu}$ is the proximal operator of the regularization function $g(Z)$. We can also consider $g(Z)$ as a regularization term for the shared variable Z , and several regularization patterns can be used in the aggregation step (12). Finally, these brackets $\langle \cdot \rangle$ denotes the mean of the corresponding variable over all the domains.

4.2 Proximal Operator Options

The proximal operator [24] for a function $g(X)$ is given by the formula:

$$\text{prox}_{g, \mu}(Z) = \arg \min_X g(X) + \frac{1}{2\mu} \|X - Z\|_F^2. \quad (14)$$

This operator solves a double-target optimization problem. It minimizes the function $g(X)$ while trying to keep X as close as possible to some value Z , and the parameter μ controls this trade-off. In this paper, we consider two simple cases of proximal operators to use in the step (12) of the ADMM iterations, and here we list these options:

4.2.1 Identity Operator. For the case of $g(X) = 0$, the proximal operator will be:

$$\text{prox}_{0, \mu}(Z) = Z. \quad (15)$$

This is a straightforward option that was used in many previous works [25, 32, 35].

4.2.2 L_2 Regularization. For $g(X) = \frac{\lambda}{2} \|X\|_F^2$, which is the squared Frobenius norm, the proximal operator is:

$$\text{prox}_{\|\cdot\|_F^2, \mu}(Z) = \arg \min_X \frac{\lambda}{2} \|X\|_F^2 + \frac{1}{2\mu} \|X - Z\|_F^2. \quad (16)$$

Solving this unconstrained problem is easy, $\frac{\partial f}{\partial X} = \lambda X + \frac{1}{\mu}(X - Z) = 0$. Then,

$$\text{prox}_{\|\cdot\|_F^2, \mu}(Z) = \frac{\mu}{1 + \lambda\mu} Z. \quad (17)$$

However, more options can be used in this step, like projection to a specific set of matrices, or volume maximization. These options are left for future research.

4.3 Details of the used ALS solver

We adapt the solver proposed by Rendle et al. [27] to solve the optimization problem in step 4. of the Algorithm (1). We chose that notation of ALS since it helps to guide the hyperparameter search iteratively rather than jointly. Rendle et al. used the following objective function for the ALS problem:

$$L(X, Y) = L_S(X, Y) + L_I(X, Y) + R(X, Y), \quad (18)$$

$$L_S(X, Y) = \sum_{(u,i) \in S} (x_u^T y_i - 1)^2, \quad (19)$$

$$L_I(X, Y) = \alpha \sum_{u \in U} \sum_{i \in I} (x_u^T y_i)^2, \quad (20)$$

$$R(X, Y) = \lambda \left(\sum_{u \in U} (|\mathcal{I}_u| + \alpha |\mathcal{I}|)^v \|x_u\|^2 + \sum_{i \in I} (|\mathcal{U}_i| + \alpha |\mathcal{U}|)^v \|y_i\|^2 \right). \quad (21)$$

Where $x_u, y_i \in \mathbb{R}^d$ are the latent factors for the user u and the item i , respectively. The loss objective in (18) consists of 3 terms; the first term L_S is the squared error between the predicted score and the observed interaction for all the pairs (user/item) in the dataset S . The next term L_I is computed for all pairs in $\mathcal{U} \times \mathcal{I}$ and represents the difference of the predicted score from 0. This loss helps the model not to learn a naive predictor $x_u y_i = 1, \forall (u, i) \in \mathcal{U} \times \mathcal{I}$. The hyperparameter α controls the contribution of L_I in the total loss function. The last term is the regularization term R , which is controlled by the global regularization parameter λ .

The scheme of individual regularization weights depends on the number of observed interactions for a specific user $|\mathcal{I}_u|$ or item $|\mathcal{I}_i|$ and on the size of the users set $|\mathcal{U}|$ and items set $|\mathcal{I}|$. The scheme is controlled by the frequency-scaled regularization parameter v . This parameter takes values from the range $[0, 1]$. For $v = 0$, all user and item factors will be regularized equally. The larger v is, the more variance there will be in the regularization effect between popular and unpopular items/users.

The following substitutions are required to adapt this model to our notation (2):

$$C = \begin{cases} 1 + \alpha, & \text{if } (u, i) \in S \\ \alpha, & \text{if } (u, i) \notin S \end{cases}, \quad (22)$$

$$\Lambda_X = \text{diag}(\{|\mathcal{I}_u| + \alpha |\mathcal{I}|^v; u \in U\}), \quad (23)$$

$$\Lambda_Y = \text{diag}(\{|\mathcal{U}_i| + \alpha |\mathcal{U}|^v; i \in I\}). \quad (24)$$

4.4 Derivation of equations for the iterations of the ALS solver

We will derive the solution of a general form that can be adapted to solve multiple steps of the CDIMF algorithm with proper substitutions. The generic formula is:

$$A = \arg \min_A \frac{1}{2} \|C \circ (P - AB^T)\|_F^2 + \frac{\lambda}{2} \|RA\|_F^2 + \frac{\rho}{2} \|A - H\|_2^2. \quad (25)$$

This problem is convex with respect to A , so the optimal A is the solution of $\partial f / \partial A = 0$, then:

$$\frac{\partial f}{\partial A} = - \left[C^{\circ 2} \circ (P - AB^T) \right] B + \lambda R^2 A + \rho(A - H) = 0, \quad (26)$$

where $C^{\circ 2}$ is the element-wise square. We can solve the problem for a single column in A , let us denote a, h for each row of matrices A, H . We also use p, c for a row of P, C respectively, and r is a scalar since R is a diagonal matrix.

$$- \left[C^{\circ 2} \circ (p - a_{row} B^T) \right] B + \lambda r^2 a_{row} + \rho(a_{row} - h_{row}) = 0,$$

$$- \left[(p_{row} - a_{row} B^T) \circ c_{row}^{\circ 2} \right] B + \lambda r^2 a_{row} + \rho(a_{row} - h_{row}) = 0,$$

$$\left[(a_{row} B^T) \circ c_{row}^{\circ 2} \right] B + \lambda r^2 a_{row} + \rho a_{row} = (p_{row} \circ c_{row}^{\circ 2}) B + \rho h_{row},$$

$$a_{row} B^T \text{diag}(c_{row}^{\circ 2}) B + \lambda r^2 a_{row} + \rho a_{row} = (p_{row} \circ c_{row}^{\circ 2}) B + \rho h_{row}.$$

Transpose:

$$\left(B^T \text{diag}(c^{\circ 2}) B + (\rho + \lambda r^2) I_d \right) a = B^T (p \circ c^{\circ 2}) + \rho h. \quad (27)$$

Solving this linear system updates a single user/item factor. Note that the problem (27) for each factor is independent of the other factors, and this feature allows great possibility for parallelism. With a very similar derivation, we can get a similar formula to update B . So, here we list substitutions for solving different problems in the steps of the algorithm (1):

- For shared users: $A = X, B = Y, H = Z - U$
- For non-shared users: $A = X, B = Y, H = 0$
- For items: $A = Y, B = X, H = 0$ and p, c will represent the corresponding *columns* of the matrices P, C

We remind readers that X are local users' factors, Y are items' factors, Z are global users' factors, and U are dual variables related to the equality constraints.

4.5 Algorithm and Complexity

The pseudocode of the algorithm running on each side (domain) is presented in Algorithm (1). In this algorithm, the matrices of user and item factors are initialized with random values from a normal distribution of standard deviation σ . This hyperparameter is recommended to be normalized by \sqrt{d} , where d is the embedding size of the factors. It helps to control the values of the dot products of factors at the beginning of training [27].

The complexity of one epoch for the step of updating local variables is the same as for standard ALS for a sparse matrix. It is $O(d(|S| + d^2|\mathcal{U}| + d^2|\mathcal{I}|))$, where $|\mathcal{S}|, |\mathcal{U}|, |\mathcal{I}|$ are the size of observed training set, the number of users, and the number of items, respectively. For the proposed options for the proximal operator, the update of the global variable Z is just the average of $X_i + U_i$ multiplied by a scalar. Thus, its complexity is $O(Nd|\mathcal{U}|)$, where N is the number of domains.

As illustrated in Section 4.4, the task of updating individual user/item factors is *embarrassingly parallel*. In other words, each column of factors' matrices X_i, Y_i can be computed in parallel without dependency on the others. Exploiting this feature in implementation in addition to the sparse nature of the interaction matrices can achieve significant speedup in practice [29].

Algorithm 1: CDIMF psuedocode for each domain i

```

init:  $X_i, Y_i \sim \mathcal{N}(0, \sigma/\sqrt{d})$ 
init:  $U_i = 0, Z = 0$ 
for  $k \leftarrow 1$  to  $K$  do
     $X_i^{k+1} \leftarrow \arg \min_{X_i} \|C_i \circ (P_i - X_i Y_i^{kT})\|_F^2 + \|\Lambda_{X_i} X_i\|_F^2$ 
     $\quad + \frac{\rho}{2} \cdot \|X_i - Z^k + U_i^k\|_2^2$ 
     $Y_i^{k+1} \leftarrow \arg \min_{Y_i} \|C_i \circ (P_i - X_i^{k+1} Y_i^T)\|_F^2 + \|\Lambda_{Y_i} Y_i\|_F^2$ 
    send:  $X_i^{k+1} + U_i^k$ 
    recieve:  $X_j^{k+1} + U_j^k$  for  $j \in [1, 2, \dots, N] \setminus \{i\}$ 
     $\langle X \rangle^{k+1} \leftarrow \frac{1}{N} \sum_{i=1}^N X_i^{k+1}$ 
     $\langle U \rangle^k \leftarrow \frac{1}{N} \sum_{i=1}^N U_i^k$ 
     $Z^{k+1} \leftarrow \text{prox}_{g, \frac{1}{\rho N}} (\langle X \rangle^{k+1} + \langle U \rangle^k)$ 
     $U_i^{k+1} \leftarrow U_i^k + (X_i^{k+1} - Z^{k+1})$ 
return  $X_i, Y_i$ 

```

5 EXPERIMENTAL SETUP

In this section, we illustrate the methodology we followed in our experiments. First, we provide details about the datasets used and the preprocessing procedure. Then, we describe the evaluation method and metrics, and finally, we list the baselines we used to compare our model with. The code to reproduce experiments is provided in the corresponding GitHub repository¹.

5.1 Datasets

Following many previous CDR works, we conduct experiments under the public Amazon dataset² for a fair comparison. We used preprocessed versions of these datasets provided by [3], and these datasets were also used by [4, 5, 20]. However, we provided a preprocessing pipeline for further experiments.

In our experiments, we tested our models in two cross-domain scenarios: the dual-intra-users and the dual-inter-users. For the first scenario, we aim to increase recommendation quality for users inside the domain (*warm-start* problem). We first filtered out unshared users to create the dataset train/test splits. Then, one item is sampled from each user history to form the test set (this method is known as *leave-one-out*) as shown in Figure 1 (a). In the second scenario, we aim to recommend items for users "migrating" from one domain to another. It is a *cold-start* problem from the point of view of the destination domain, as users do not have a history there. To emulate this scenario, we sampled the complete history of two disjoint sets of shared users from each domain, as shown in Figure 1 (b). In the cold start scenario, note that test users in one domain are considered unshared in the other since their history is

Table 1: Dataset statistics

Scenario	Dataset	Train set				Test set		
		#users	#items	#ratings	#shared	#users	#items	#ratings
intra-domain (warm-start)	Sport	9928	30796	92612	9928	8326	5734	8326
	Cloth	9928	39008	87829		7540	5495	7540
	Phone	3325	17709	50407	3325	2559	1795	2559
	Electronics	3325	38706	115554		2560	2099	2559
inter-domain (cold-start)	Sport	26346	12655	163291	25356	897	2698	3546
	Cloth	40839	17943	187880		907	2519	3085
	Video	19240	8751	156091	19014	216	1157	1458
	Games	24799	12319	155036		212	1138	1304

sampled out. Thus, the set of users with the shared factors will be smaller than the actual intersection of users of datasets.

As mentioned in [20], the density of the preprocessed data was increased by filtering out users of less than 5 interactions and items with less than 10 interactions. For the warm start scenario, there are two pairs of datasets: **Sport & Cloth** and **Phone & Electronics**. There are also two pairs of datasets of the cold start case. These datasets are **Sport & Cloth** (with different preprocessing and splitting patterns) and **Games & Video**. Some statistics about the size of datasets are shown in Table 1.

5.2 Evaluation Method and Metrics

Sampled metrics are getting increasingly popular in many research works. They provide faster computations as relevant items are ranked against a random sample of negative items, not the entire catalog. However, they may be biased when comparing different models [17]. However, with a sufficiently large sample size, we claim to have an unbiased estimation of model performance. In the following experiments, we sampled 999 negative items for each item in the test set. Then, we rank this item list and evaluate the top 10 recommended items with the two famous metrics: *NDCCG* (Normalized Discounted Cumulative Gain) and *HR* (Hit Ratio).

We used a separate validation set to tune the hyperparameters. In the preprocessed datasets provided by [3], there are separate files for validation sets for each domain in the warm-start scenario, but not for the cold-start. Thus, for the latter scenario, we split the provided training set into training and validation sets. We also used the *leave_one_out* method to make this validation set, where one item is sampled out of each user history. For final testing, we used the provided test sets.

5.3 Baselines

Single-domain baselines: We first compare with BPRMF [26], NeuMF [10]. These are typical baseline methods for industry because of the straightforward idea and simple architecture. The latest graph-based methods were included in the comparison: NGCF [33] and LightGCN [9]. These models aggregate the high-order neighbor information for the user/item. For single-domain methods, all domain interaction data were mixed as one domain to train them.

However, CDIMF is a cross-domain variant of ALS [27]. Thus in both scenarios, we report the results of a single ALS model trained on the joined datasets. We joined the training dataframes so the resulting dataframe has a number of users equal to the size of the union of user sets and a number of items equal to the sum of the numbers of items in both domains since the item sets are disjoint.

¹<https://github.com/aa-samra/cd-imf>

²https://jmcauley.ucsd.edu/data/amazon/index_2014.html

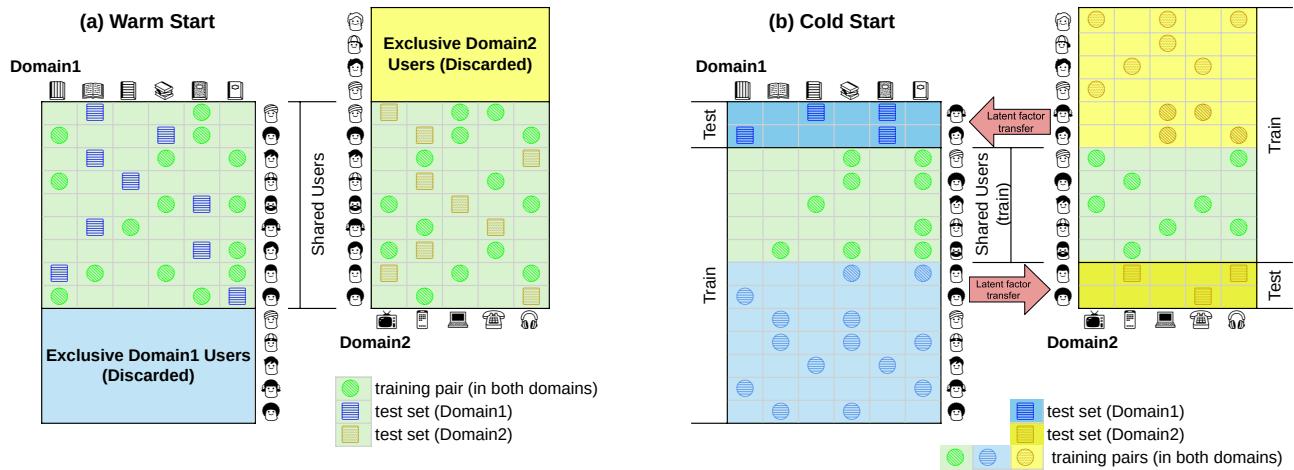


Figure 1: Train/test data splits for (a) intra-domain recommendation scheme (warm start) and (b) inter-domain recommendation scheme (cold start); circles refer to training interactions, and squares refer to test interactions. Arrows refer to latent factors transfer for testing

This baseline aims to show the performance of full cross-domain sharing with no privacy restrictions.

For the warm-start specifically, we report another baseline. It is the ALS trained on each domain separately. This setting is equivalent to running the CDIMF code with sharing parameter $\rho = 0$. It would help understand whether sharing knowledge across domains has a positive or negative impact on recommendation quality compared to the single domain case. In the cold start, there is no point in comparing with ALS on separate datasets since user factors on each domain would be learned independently, and transferring these factors to other domains will result in systematically wrong recommendations.

Intra-domain CDR baselines: For this scenario, we compare with MLP-based models like CoNet [12] and DDTCDR [19]. These models assign an encoder for each domain and learn different dual transferring modules between source and target domains. In addition, recent GNN-based models like Bi-TGCF [20] and DisenCDR [4] are included in the comparison. We also compare with UniCDR [3], a model that follows a contrastive learning approach and scores good metrics.

Inter-domain CDR baselines: For this category of CDR methods, the models SSCDR [15], TMCDD [41], and SA-VAE [28] mainly rely on the learning embedding and then mapping them. They train an alignment function to project the user embeddings from the source domain to the target domain. Besides, CDRIB [5] depends on the variational information bottleneck principle to learn an unbiased representation to encode the domain invariant information. We also compare with UniCDR [3] as it is a model designed to work in both cold-start and warm-start settings.

6 RESULTS

We report the results of our experiments in this section. The first set of experiments aims to compare the proposed model with the

baselines, and the second group of experiments explores the behavior of CDIMF and the effect of different hyperparameters on it.

6.1 Performance Comparison

Table 2 reports the performance metrics of the proposed model and the compared models for warm start. CDIMF demonstrates significant improvement in most experiments for this scenario. However, these advancements are more noticeable in some experiments than in others. Although CDIMF fails to score the top HR@10 on some datasets, it scores the best NDCG@10. The latter metric is more critical and less noisy than the Hit Rate since it provides a more comprehensive evaluation of the overall quality of recommendations and how well relevant items are positioned in the list. The results of the ALS model trained on separate datasets demonstrate that sharing cross-domain knowledge does not guarantee better performance in some experiments. One example is in Table 2, where ALS achieved the highest metrics for the **Phone** domain in the second experiment.

Table 3 shows the performance metrics of the CDIMF and the compared models for the cold start scenario. CDIMF shows a significant increase in metrics in most experiments in this scenario as well. However, these improvements are less significant than in warm-start experiments. We also can notice that the ALS on the joined dataset score less than the top performing baselines in the **Game-Video** experiment, even less than the CDIMF, unlike other experiments. That may reflect that this pair of datasets in particular contains patterns that can't be learned using models based on simple matrix factorization.

One general observation that can be concluded from both tables is that CDIMF scores are slightly lower than the ALS model trained on the mixed dataset, e.g., incorporated data from both domains into a single dataset. *These correlated scores support our claim that*

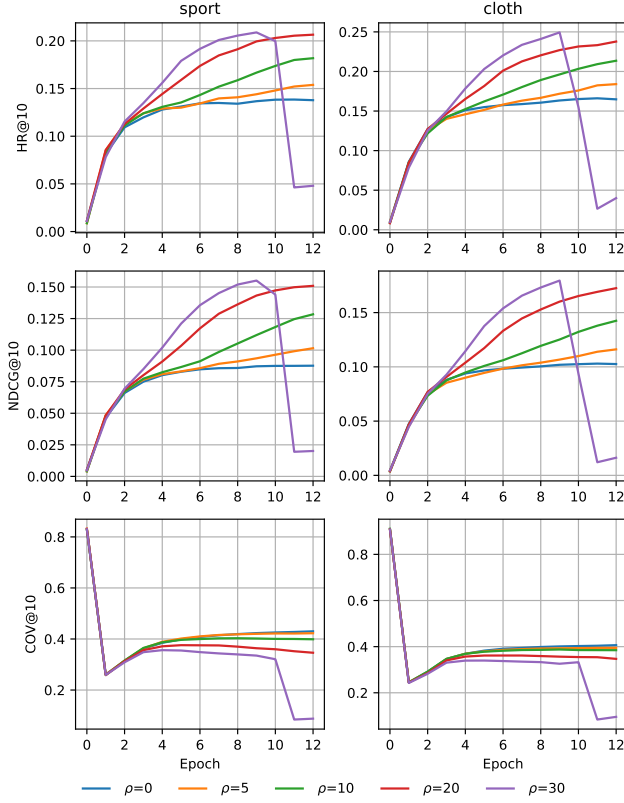


Figure 2: HR@10, NDCG@10, and coverage changes in the warm start problem for several values of sharing parameter ρ . For $\rho = 0$, a single domain ALS is running on each side.

CDIMF is a cross-domain variant of ALS. Unlike ALS on the joint datasets, CDIMF does not share explicit user data between domains.

The last point to mention is that basic single-domain methods trained on each domain dataset separately show competing performance to cross-domain models in some experiments. This observation raises the question of whether sharing cross-domain knowledge is always useful for all pairs of domains. This question may establish an interesting research topic about the "compatibility" of datasets and how to estimate the expected benefit of cross-domain knowledge sharing.

6.2 Hyperparameters Discussion

Apart from the hyperparameters of the ALS solver described in Section 4.3, CDIMF has three characteristic hyperparameters: the sharing (penalty) parameter, the proximal operator, and the aggregation period. We conducted several experiments to understand the effect of each hyperparameter on the behavior of CDIMF.

6.2.1 The effect of sharing parameter ρ . In Figure 2, we can see that the sharing factor ρ plays a role in the learning rate of CDIMF as it controls the learning speed. For $\rho = 0$ the iteration in (11) would

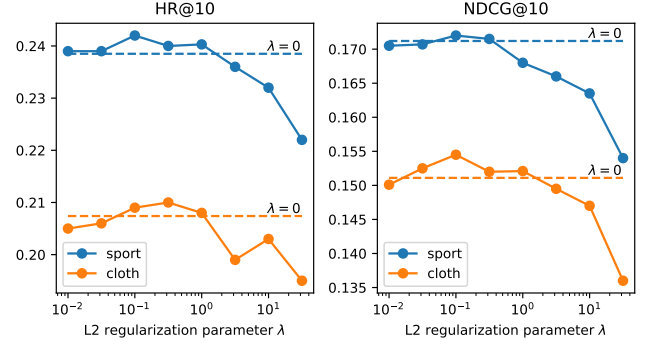


Figure 3: The effect of L_2 regularization on performance of CDIMF. Scenario: warm start, dataset: Cloth&Sport

be:

$$X_i^*, Y_i^* = \arg \min_{X_i, Y_i} \|C_i \circ (P_i - X_i Y_i^T)\|_F^2 + \|\Lambda_{X_i} X_i\|_F^2 + \|\Lambda_{Y_i} Y_i\|_F^2,$$

which is simply a single-domain ALS. The model learns cross-domain knowledge more rapidly as the sharing parameter increases. The model starts to diverge for higher values of ρ , leading to the learning of divergent item factors. This results in a few items with the largest norms, which score larger dot products with any user factor. Therefore, these items are usually recommended. This explains why the coverage drops dramatically to a near-zero value when divergence occurs.

6.2.2 The effect of proximal operator. In this paper, we considered two simple options for the proximal operator in (12); the identity function and the L_2 regularization. Figure 3 demonstrates the effect of L_2 regularization on the CDIMF models. Although choosing the optimal value for the regularization parameter λ can enhance the performance slightly over the case of no regularizer, reported changes are relatively small. Finding the best value of λ increases the time of hyperparameter tuning without significant improvement. Thus, all numbers in Tables 2 and 3 are for the CDIMF with identity operator.

6.2.3 The effect of aggregation period. The typical ADMM iterations assume having the optimal solution locally (11) before updating the common variable (12), experiments show that sharing a sub-optimal solution (resulting from a single ALS epoch) frequently is better. However, the communication load would increase in this situation. Figure 4 shows that the more frequent aggregations are, the faster performance growth is. This finding may allow asynchronous updates of the global variable. Supposing that the global variable Z was not updated in some epochs due to communication issues, a local ALS iteration could be run using the old version of Z and still achieve decent, yet lower, improvement.

6.3 Privacy Considerations

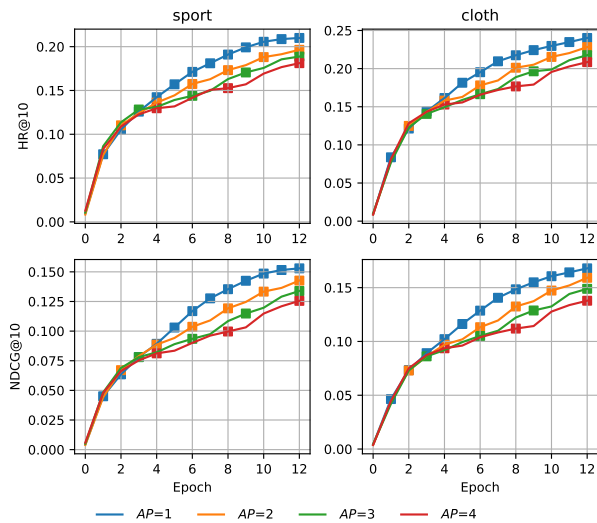
The variable shared by the node in each domain is $X_i + U_i$, which is a mixture of the local factors of the users and the dual variable related to the equality constraint $X_i - Z = 0$. We hypothesize that the explicit user history in a particular domain cannot be extracted

Table 2: Performance comparison of our proposed model on dual-intra-domain (warm start) recommendations. The asterisk sign indicates that the single-domain method trained separately on local data outperformed all cross-domain methods.

Datasets		Sport		Cloth		Electronic		Phone	
Metrics@10		HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Single-domain methods	BPRMF	10.43	5.41	11.53	6.25	15.71	9.19	16.32	8.53
	NeuMF	10.74	5.46	11.18	6.02	16.17	9.24	15.84	8.02
	LightGCN	13.19	6.94	13.58	7.29	19.17	10.28	23.25	12.72
	ALS (separate)	16.67	10.43	13.9	8.86	22.11	14.01	29.54*	18.74*
	ALS (joined)	25.69	18.44	22.15	16.4	21.37	14.91	32.51	22.64
Cross-domain methods	CoNet	12.09	6.41	12.40	6.62	17.22	9.86	17.66	9.30
	DDTCDR	11.86	6.37	12.54	7.13	18.47	11.08	17.23	8.58
	Bi-TGCF	14.83	7.95	14.68	7.93	22.14	12.20	25.71	13.93
	DisenCDR	17.55	9.46	16.31	9.03	24.57	14.51	28.76	16.13
	UniCDR	<u>18.37</u>	<u>10.98</u>	<u>17.85</u>	<u>11.20</u>	<u>22.92</u>	13.83	24.72	13.77
Our method	CDIMF	23.85	17.12	20.74	15.11	22.66	14.65	<u>28.57</u>	18.3

Table 3: Performance comparison of our proposed model on dual-inter-domain (cold start) recommendations

Datasets		Sport		Cloth		Game		Video	
Metrics@10		HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Single-domain methods	BPRMF	5.75	3.16	6.75	3.26	3.77	1.89	4.46	2.36
	NGCF	7.22	3.63	7.07	3.48	5.14	2.73	7.41	3.87
	ALS (joined)	15.01	10.61	13.68	9.36	7.44	3.83	8.16	4.64
Cross-domain methods	SSCDR(CML)	7.27	3.75	6.12	3.06	3.48	1.59	5.51	2.61
	TMCDR	7.18	3.84	8.11	5.05	5.36	2.58	8.85	4.41
	SA-VAE	7.51	3.72	7.21	4.59	5.84	2.78	7.46	3.71
	CDRIB	<u>12.04</u>	6.22	12.19	6.81	8.51	4.58	13.17	6.49
	UniCDR	<u>11.20</u>	<u>7.04</u>	<u>12.48</u>	<u>7.52</u>	8.78	4.63	<u>10.74</u>	5.89
Our method	CDIMF	14.52	9.51	13.35	9.45	9.19	4.53	8.50	4.36

**Figure 4: HR@10 and NDCG@10 changes in the warm start problem for several values of aggregation period (AP), squares refer to epochs of shared variable update**

unless the item factors of that domain are exposed. Thus, communication could be held over an insecure channel as long as the

domains' servers are well protected. However, further investigation of the possible reconstruction attacks is needed in this aspect.

7 CONCLUSION

In this paper, we introduced CDIMF, a model that upgrades the conventional ALS method to the cross-domain context. Our experimental results demonstrate that CDIMF outperforms other state-of-the-art CDR methods in most cases, often surpassing even more elaborate neural network-based approaches. Consequently, among the key advantages of our CDIMF approach are competitive performance and superior computational efficiency inherited from ALS, which makes it plausible for practical use cases.

CDIMF model has been tested in the dual domain case for both warm-start and cold-start problems. Moreover, the extension to the multi-domain case can be achieved with minor adjustments. The task here is to deal with cases of users who have a history in more than one, but not all, domains. Further experiments should be conducted to evaluate possible solutions for this non-homogenous setting. The CDIMF model also has the advantage of notation symmetry, and it can extend to item-overlapping scenarios by swapping the notations.

REFERENCES

- [1] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2010. Distributed optimization and statistical learning via the alternating direction method of multipliers. , 122 pages. Issue 1. <https://doi.org/10.1561/2200000016>

- [2] Iván Cantador, Ignacio Fernández-Tobías, Shlomo Berkovsky, and Paolo Cremonesi. 2015. *Cross-domain recommender systems*. Springer US, 919–959. https://doi.org/10.1007/978-1-4899-7637-6_27
- [3] Jiangxia Cao, Shaoshuai Li, Bowen Yu, Xiaobo Guo, Tingwen Liu, and Bin Wang. 2023. Towards Universal Cross-Domain Recommendation. *WSDM 2023 - Proceedings of the 16th ACM International Conference on Web Search and Data Mining*, 78–86. <https://doi.org/10.1145/3539597.3570366>
- [4] Jiangxia Cao, Xixun Lin, Xin Cong, Jing Ya, Tingwen Liu, and Bin Wang. 2022. Disencdr: Learning disentangled representations for cross-domain recommendation. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*. 267–277.
- [5] Jiangxia Cao, Jiawei Sheng, Xin Cong, Tingwen Liu, and Bin Wang. 2022. Cross-domain recommendation to cold-start users via variational information bottleneck. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2209–2223.
- [6] Jiawei Chen, Can Wang, Sheng Zhou, Qihao Shi, Jingbang Chen, Yan Feng, and Chun Chen. 2020. Fast Adaptively Weighted Matrix Factorization for Recommendation with Implicit Feedback. *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (04 2020), 3470–3477. <https://doi.org/10.1609/aaai.v34i04.5751>
- [7] Ben Frederickson et al. 2023. *Fast Python Collaborative Filtering for Implicit Datasets*. <https://github.com/benfred/implicit>
- [8] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems (Copenhagen, Denmark) (RecSys '19)*. Association for Computing Machinery, New York, NY, USA, 101–109. <https://doi.org/10.1145/3298689.3347058>
- [9] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [10] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [11] Xiangnan He, Hanwang Zhang, Min Yen Kan, and Tat Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. *SIGIR 2016 - Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 549–558. <https://doi.org/10.1145/2911451.2911489>
- [12] Guangmeng Hu, Yu Zhang, and Qiang Yang. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 667–676.
- [13] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets.
- [14] Nouhaila Idrissi and Ahmed Zellou. 2020. A systematic literature review of sparsity issues in recommender systems. *Social Network Analysis and Mining* 10 (12 2020), Issue 1. <https://doi.org/10.1007/s13278-020-0626-2>
- [15] SeongKu Kang, Junyong Hwang, Dongha Lee, and Hwanjo Yu. 2019. Semi-supervised learning for cross-domain recommendation to cold-start users. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1563–1572.
- [16] Hyeyoung Kim, Jae-Nam Lee, and Jaemin Han. 2010. The role of IT in business ecosystems. *Commun. ACM* 53, 5 (may 2010), 151–156. <https://doi.org/10.1145/1735223.1735260>
- [17] Walid Krichene and Steffen Rendle. 2020. On sampled metrics for item recommendation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 1748–1757.
- [18] Bin Li, Qiang Yang, and Xiangyang Xue. 2009. Can movies and books collaborate? cross-domain collaborative filtering for sparsity reduction. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (Pasadena, California, USA) (IJCAI'09)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2052–2057.
- [19] Pan Li and Alexander Tuzhilin. 2019. DDTCDR: Deep Dual Transfer Cross Domain Recommendation. [arXiv:1910.05189 \[cs.LG\]](https://arxiv.org/abs/1910.05189)
- [20] Meng Liu, Jianjun Li, Guohui Li, and Peng Pan. 2020. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 885–894.
- [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, IJCAI-17*. 2464–2470. <https://doi.org/10.24963/ijcai.2017/343>
- [22] Xia Ning and George Karypis. 2011. SLIM: Sparse Linear Methods for Top-N Recommender Systems. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining (ICDM '11)*. IEEE Computer Society, USA, 497–506. <https://doi.org/10.1109/ICDM.2011.134>
- [23] Weike Pan, Evan W. Xiang, Nathan N. Liu, and Qiang Yang. 2010. Transfer learning in collaborative filtering for sparsity reduction. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (Atlanta, Georgia) (AAAI'10)*. AAAI Press, 230–235.
- [24] Neal Parikh, Stephen Boyd, et al. 2014. Proximal algorithms. *Foundations and trends® in Optimization* 1, 3 (2014), 127–239.
- [25] Tran-Dinh Quoc, Nhan H Pham, Dzung T Phan, and Lam M Nguyen. 2021. FedDR-Randomized Douglas-Rachford Splitting Algorithms for Nonconvex Federated Composite Optimization.
- [26] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [27] Steffen Rendle, Walid Krichene, Li Zhang, and Yehuda Koren. 2021. Revisiting the Performance of iALS on Item Recommendation Benchmarks. (10 2021). <http://arxiv.org/abs/2110.14037>
- [28] Aghiles Salah, Thanh Binh Tran, and Hady Lauw. 2021. Towards source-aligned variational models for cross-domain recommendation. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 176–186.
- [29] Abdulaziz Samra. 2024. *Comparison of Several Implementations for iALS speedup*. <https://doi.org/10.5281/zenodo.11085635>
- [30] 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 (Las Vegas, Nevada, USA) (KDD '08)*. Association for Computing Machinery, New York, NY, USA, 650–658. <https://doi.org/10.1145/1401890.1401969>
- [31] Harald Steck, Maria Dimakopoulou, Nickolai Riabov, and Tony Jebara. 2020. ADMM SLIM: Sparse Recommendations for Many Users. In *Proceedings of the 13th International Conference on Web Search and Data Mining (Houston, TX, USA) (WSDM '20)*. Association for Computing Machinery, New York, NY, USA, 555–563. <https://doi.org/10.1145/3336191.3371774>
- [32] Han Wang, Siddhartha Marella, and James Anderson. 2022. FedADMM: A Federated Primal-Dual Algorithm Allowing Partial Participation. (3 2022). <http://arxiv.org/abs/2203.15104>
- [33] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*. 165–174.
- [34] Erkang Xue, Ruixin Guo, Feng Zhang, Lizhe Wang, Xiaohan Zhang, and Guangzhi Qu. 2019. Distributed Differentially Private Matrix Factorization Based on ADMM. In *2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. 2502–2507. <https://doi.org/10.1109/HPCC/SmartCity/DSS.2019.00350>
- [35] Wen Yan, Chunguo Li, Yongming Huang, and Luxi Yang. 2019. Low-Complexity Decentralized Recommendation System with Similarity Constraints. *IEEE Access* 7 (2019), 146922–146938. <https://doi.org/10.1109/ACCESS.2019.2946485>
- [36] Zhi-Qin Yu, Xing-Jian Shi, Ling Yan, and Wu-Jun Li. 2014. Distributed Stochastic ADMM for Matrix Factorization. In *Proceedings of the 23rd ACM International Conference on Information and Knowledge Management (Shanghai, China) (CIKM '14)*. Association for Computing Machinery, New York, NY, USA, 1259–1268. <https://doi.org/10.1145/2661829.2661996>
- [37] Tianzi Zang, Yanmin Zhu, Haobing Liu, Ruohan Zhang, and Jiadi Yu. 2022. A Survey on Cross-domain Recommendation: Taxonomies, Methods, and Future Directions. *ACM Trans. Inf. Syst.* 41, 2, Article 42 (dec 2022), 39 pages. <https://doi.org/10.1145/3548455>
- [38] Feng Zhang, Erkang Xue, Ruixin Guo, Guangzhi Qu, Gansen Zhao, and Albert Y. Zomaya. 2022. DS-ADMM++: A Novel Distributed Quantized ADMM to Speed up Differentially Private Matrix Factorization. *IEEE Transactions on Parallel and Distributed Systems* 33, 6 (2022), 1289–1302. <https://doi.org/10.1109/TPDS.2021.3110104>
- [39] Zhi-Lin Zhao, Ling Huang, Chang-Dong Wang, and Dong Huang. 2018. Low-Rank and Sparse Cross-Domain Recommendation Algorithm. In *Database Systems for Advanced Applications*, Jian Pei, Yannis Manolopoulos, Shazia Sadiq, and Jianxin Li (Eds.). Springer International Publishing, Cham, 150–157.
- [40] Feng Zhu, Yan Wang, Chaochao Chen, Jun Zhou, Longfei Li, and Guanfang Liu. 2021. Cross-Domain Recommendation: Challenges, Progress, and Prospects. (3 2021). <http://arxiv.org/abs/2103.01696>
- [41] Yongchun Zhu, Kaikai Ge, Fuzhen Zhuang, Ruobing Xie, Dongbo Xi, Xu Zhang, Leyu Lin, and Qing He. 2021. Transfer-meta framework for cross-domain recommendation to cold-start users. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 1813–1817.