PAT: Preference-Aware Transfer Learning for Recommendation with Heterogeneous Feedback

Feng Liang, Wei Dai, Yunfeng Huang, Weike Pan* and Zhong Ming*
National Engineering Laboratory for Big Data System Computing Technology, Shenzhen University
Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), Shenzhen University
College of Computer Science and Software Engineering, Shenzhen University
Shenzhen, China
{liangfeng2018,daiwei20171,huangyunfeng2017}@email.szu.edu.cn, {panweike,mingz}@szu.edu.cn

***Abstract*—In this paper, we study an important recommendation problem with heterogeneous feedback of users' grade scores such as 5-star grade scores and like/dislike binary ratings assigned to items. As a response, we address the problem from a transfer learning perspective, i.e., taking the grade scores as the target data and the binary ratings as the auxiliary data, in order to share knowledge between two different types of data more sufficiently. Technically, besides the observed explicit feedback of grade scores and binary ratings, we propose to exploit the implicit preference context beneath the feedback, which is incorporated into the prediction process of users' grade scores to items. Finally, we develop a novel and generic transfer learning solution, i.e., preference-aware transfer (PAT), which embodies several recent algorithms as parts of our solution and special cases. To verify our novel solution, we then conduct extensive empirical studies on two large and public datasets and find that our PAT performs significantly better than the state-of-the-art methods.**

***Index Terms*—Transfer Learning, Implicit Preference Context, Matrix Factorization, Collaborative Filtering**

I. Introduction

Recommender systems are now ubiquitous and benefit us in many aspects of our daily lives. For example, the news article [1] recommended by a news website is exactly what we follow closely, the clothing [2] recommended by a dress shopping system is what we like exactly, and the event [3] recommended by a mobile recommendation system is what we interest currently. With all kinds of recommendation services, people can easily obtain interesting information without wasting too much time in searching.

An effective recommendation algorithm is the cornerstone to a successful recommender system for the purpose of offering users sensible suggestions. Collaborative filtering [4], [5] is such a commonly used technique that has aroused considerable research interest for decades because of its effectiveness and scalability. One of the tasks of collaborative filtering (CF) is to predict accurate grade scores to items that users never rated in the past, which is also known as rating prediction. Among various CF methods for predicting ratings, one of the most popular methods is matrix factorization (MF) [6] in which probabilistic matrix factorization (PMF) [7] is a remarkable model widely applied in industry due to its efficiency and effectiveness. Given a (user, item) matrix where each element is a score that a user assigns to an item, PMF [7] first decomposes the (user, item) matrix into two low-rank matrices representing users' latent features and items' latent features, respectively, then calculates the inner product of a user's feature vector and an item's feature vector as the predicted score for the corresponding (user, item) pair. As a classical CF method, PMF [7] has exhibited its success in learning a recommendation model from homogeneous explicit feedback, i.e., the grade scores only. However, most real applications are confronted with the problem of data sparsity, which means that a large number of elements in the (user, item) matrix are unobserved. Therefore, it is necessary to consider other types of available explicit feedback into modeling in a bid to improve the recommendation accuracy, which is the problem of collaborative recommendation with heterogeneous explicit feedback studied in this paper.

Collaborative recommendation with heterogeneous explicit feedback has been well studied in [8], where a unified transfer learning [9] framework, i.e., transfer by mixed factorization (TMF) [8] is proposed based on PMF [7] in an attempt to improve the performance by transferring knowledge between grade scores and like/dislike binary ratings as much as possible. Experiments show that these two types of data are complementary for the prediction accuracy.

Besides introducing additional explicit binary feedback, there are some implicit preference context hidden in grade scores ignored by TMF [8], such as one-class preference context (OPC) and multiclass preference context (MPC), which are included in SVD++ [10] and MF-MPC [11], respectively. Both the two methods reach a higher accuracy than the corresponding former versions.

In this paper, we discuss several methods extended from PMF [7] that integrate heterogeneous explicit feedback and implicit preference context independently in a proper way. Motivated by those prominent works, we take both heterogeneous explicit feedback and implicit preference context into consideration, and then propose a novel and generic framework, i.e., preference-aware transfer (PAT), which absorbs PMF [7], SVD++ [10], TMF [8], and MF-MPC [11], and can also be reduced to any of them with proper configurations.

We organize the rest of the paper as follows. In Section 2, we discuss some closely related works. In Section 3, we

*:Corresponding author.

introduce the proposed framework in detail. In Section 4, we conduct extensive empirical studies. Finally, in Section 5, we give a brief summary and future directions.

II. Related Work

A. Probabilistic Matrix Factorization

Probabilistic matrix factorization (PMF) [7] is a dominant recommendation model that takes the explicit grade score matrix as input and outputs the learned low-rank feature vectors of users and items. At the stage of prediction, the inner product of a user's feature vector and an item's feature vector is the predicted score given by PMF. The objective of PMF can be formulated as a regression-oriented optimization problem [7].

PMF plays a vital role in dealing with a large and dense data. But in the case of lacking explicit grade records, PMF may not be a good choice for making appropriate recommendations.

B. Transfer by Collective Factorization

In practice, explicit binary ratings are more common and easier to be collected than grade scores. To overcome the problem of data sparsity, like/dislike data is considered in the framework of transfer by collective factorization (TCF) [12]. Unlike PMF that factorizes one (user, item) matrix only, with the introduction of the auxiliary data, TCF models users' personalized preferences from both grade scores and binary ratings collectively by sharing users' features and items' features. Moreover, there is a tradeoff parameter in TCF that can weigh the effect caused by the auxiliary data. Notice that when the auxiliary binary ratings are not considered, TCF is reduced to PMF.

C. Interaction-Rich Transfer by Collective Factorization

The way that using the auxiliary binary data in TCF is simple and straightforward, but the bridge constructed to transfer information between two types of data is not strong enough, so that some valuable knowledge can not be shared sufficiently. Consequently, iTCF [13] is proposed based on CMF [14] which exploits the rich interactions among the user-specific latent features of the target data and the auxiliary data when calculating the gradients of items in the model training stage. In addition, there is a tradeoff parameter in iTCF to measure the strength of interactions between user-specific latent features. And iTCF is more efficient than TCF because it uses stochastic gradient descent (SGD) as an optimization method while TCF uses batch gradient descent (BGD) [13]. Notice that when the interaction weight is 1, iTCF reduces to CMF.

D. Transfer by Mixed Factorization

Although iTCF brings in richer interactions between the target data and the auxiliary data in the model training stage whereas the way that using the user-specific and the item-specific latent feature information in the grade score prediction stage is the same as that in CMF. However, there is still some latent feature information from the auxiliary binary data being omitted in the grade score prediction stage. Hence, transfer by mixed factorization (TMF) [8] combines the feature vectors learned from two different types of data in a collective and integrative manner.

From an integrative term in objective function of TMF, the binary rating data for each user is divided into a positive item set and a negative item set, which contains items that the user likes and dislikes, respectively. Notice that when the like/dislike feedback of users to items are not considered in grade score prediction, TMF becomes iTCF.

E. Preference Context

In most practical recommendation applications, users' preference to items are generally quantified by specific scores in $\{0.5, 1, \ldots, 5\}$ or other ranges. In the training period of PMF, for each user, the score information of each item is used to fit the regression model independently, and in the prediction stage, a user's score to an item only depends on the features of the given user and the target item. However, in SVD++ [10], a user's estimated score to an item is related to other items that the user rated in the past, which are called preference context of the user. Furthermore, there is no difference among these rated items because whatever scores they are assigned, they are in the same set, or in other words, their effects are classified into the same class, which is a typical example of one-class preference context (OPC). When predicting the unobserved score, the introduction of OPC can provide a global preference context for the users . In MF-MPC [11], on the other side, rated items except the target one of a given user, i.e., preference context, are classified into several clusters in terms of the grade scores, which are named multi-class preference context (MPC). Intuitively, MPC is an advanced version of OPC which not only offers the global preference information of users, but also distinguishes the information with different values.

In a summary, TMF inherits the advantages of PMF, TCF and iTCF. On the one hand, although TMF takes advantage of explicit feedback with likes/dislikes of users to items in the auxiliary binary data in the prediction phase of users' grade scores to items (i.e., TMF exploits the implicit preference context of users from the auxiliary binary data in grade score prediction), but the implicit preference context of users in the midst of the target data is not exploited for grade score prediction. On the other hand, SVD++ and MF-MPC only exploit the preference context in the target data in order to model users' personalized preferences, and do not consider the binary ratings in auxiliary data as used in TMF. Hence, we propose a novel solution PAT which is a unified and generic framework to incorporate TMF with the one-class preference context in SVD++ and the multi-class preference context in MF-MPC, respectively, from a complementary view in order to further improve the grade score prediction accuracy.

III. Preference-Aware Transfer

A. Problem Definition

In this paper, we study a recommendation setting with heterogeneous explicit feedback, including users' grade scores and binary ratings to items. Specifically, for the input of the training data, we have a set of grade score records $\mathcal{R} = \{(u, i, r_{ui})\}$ with $r_{ui} \in \mathbb{G}$ as a grade score such as $\{0.5, 1, \ldots, 5\}$, and a set of binary rating records $\tilde{\mathcal{R}} = \{(u, i, \tilde{r}_{ui})\}$ with $\tilde{r}_{ui} \in \mathbb{B} = \{\text{like}, \text{dislike}\}$. Our goal is to estimate the grade score of each (user, item) pair in the test data $\mathcal{T}_E$, which will be measured via a pointwise loss. Notice that the binary rating records $\tilde{\mathcal{R}}$ are leveraged to improve the modeling process of the grade score records $\mathcal{R}$. We list some commonly used notations in Table I for quick index and reference.

B. Implicit Preference Context

In order to jointly model two different types of explicit feedback, i.e., r_{ui} and $\tilde{r}_{ui}$, a state-of-the-art method is proposed to approximate the grade score and binary rating simultaneously by sharing some latent variable [14],

$$\begin{cases} \hat{r}_{ui} = U_{u\cdot} V_{i\cdot}^T + b_u + b_i + \mu \\ \hat{\tilde{r}}_{ui} = W_{u\cdot} V_{i\cdot}^T \end{cases} \tag{1}$$

where the item-specific latent feature vector $V_{i\cdot}$ is shared between two factorization subtasks.

However, for the goal of grade score prediction, some implicit preference contexts are missing in the above joint modeling approach shown in (1). For example, two users u and u' with similar sets of rated items, i.e., $\mathcal{I}_u^g$ and $\mathcal{I}_{u'}^g$ with $g \in \mathbb{G}$, are likely to have similar tastes, which are called graded preference context [11]. Besides that, we may also assume that two users with similar sets of liked items (with positive feedback) or disliked items (with negative feedback) are similar in terms of their preferences [8]. Mathematically, we may represent the one-class preference context as $\bar{C}_{u\cdot}^{\mathcal{O}}$, the graded preference context as $\bar{C}_{u\cdot}^{\mathcal{G}}$, the positive preference context as $\bar{C}_{u\cdot}^{p}$, and the negative preference context as $\bar{C}_{u\cdot}^{n}$ as follows [10], [11], [8],

$$\bar{C}_{u\cdot}^{\mathcal{O}} = \delta_{\mathcal{O}} \frac{1}{\sqrt{|\mathcal{I}_u \backslash \{i\}|}} \sum_{i' \in \mathcal{I}_u \backslash \{i\}} C_{i'\cdot}^{o} \tag{2}$$

$$\bar{C}_{u\cdot}^{\mathcal{G}} = \delta_{\mathcal{G}} \sum_{g \in \mathbb{G}} \frac{1}{\sqrt{|\mathcal{I}_u^g \backslash \{i\}|}} \sum_{i' \in \mathcal{I}_u^g \backslash \{i\}} C_{i'\cdot}^{g} \tag{3}$$

$$\bar{C}_{u\cdot}^{p} = \delta_p w_p \frac{1}{\sqrt{|\mathcal{P}_u|}} \sum_{j \in \mathcal{P}_u} C_{j\cdot}^{p} \tag{4}$$

$$\bar{C}_{u\cdot}^{n} = \delta_n w_n \frac{1}{\sqrt{|\mathcal{N}_u|}} \sum_{j \in \mathcal{N}_u} C_{j\cdot}^{n} \tag{5}$$

where $\delta_{\mathcal{O}}, \delta_{\mathcal{G}}, \delta_p, \delta_n \in \{1, 0\}$ are the indicator variables, and w_p and w_n are the weights on positive feedback and negative feedback, respectively. Notice that $\bar{C}_{u\cdot}^{\mathcal{O}}$ has been used in a traditional and popular matrix factorization algorithm [10], $\bar{C}_{u\cdot}^{\mathcal{G}}$ has been used in a recent matrix factorization algorithm [11], and $\bar{C}_{u\cdot}^{p}$ and $\bar{C}_{u\cdot}^{n}$ are common tricks used in factorization-based methods [10], [8]. Our goal is to make use of all those four types of preference context in a unified framework for the studied problem in order to further improve the prediction accuracy beyond those very competitive methods.

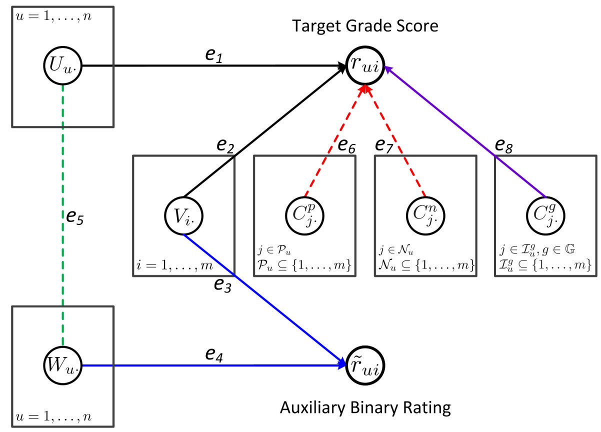

Fig. 1. Graphical model of our preference-aware transfer (PAT).

C. Transfer with Implicit Preference Context

With the preference context, we propose to incorporate them into the collective factorization framework,

$$\begin{cases} \hat{r}_{ui} = U_{u\cdot} V_{i\cdot}^T + b_u + b_i + \mu + (\bar{C}_{u\cdot}^{\mathcal{O}} + \bar{C}_{u\cdot}^{\mathcal{G}} + \bar{C}_{u\cdot}^{p} + \bar{C}_{u\cdot}^{n}) V_{i\cdot}^T \\ \hat{\tilde{r}}_{ui} = W_{u\cdot} V_{i\cdot}^T \end{cases} \tag{6}$$

which will bring two user-specific latent feature vectors of user u and user u' to be close if they have similar implicit preference context in a similar way to that of SVD++ [10]. Notice that we incorporate the preference context into the prediction rule of grade scores instead of that of binary ratings because that matches our final goal of grade score prediction rather than binary rating prediction. We illustrate the overall prediction rule in a graphical model in Fig. 1. Notice that the green dashed line is used for introducing richer interactions between the target data (i.e., grade scores) and the auxiliary data (i.e., binary ratings), and if the grade scores are classified into one cluster, then (3) will reduce to (2) and the framework of PAT in Fig. 1 will reduce to PAT-OPC simultaneously.

We then reach an objective function similar to that of CMF [14], iTCF [13] and TMF [8],

$$\min_{\Theta} \sum_{u=1}^{n} \sum_{i=1}^{m} y_{ui} f_{ui} + \lambda \sum_{u=1}^{n} \sum_{i=1}^{m} \tilde{y}_{ui} \tilde{f}_{ui} \tag{7}$$

where $f_{ui} = \frac{1}{2}(r_{ui} - \hat{r}_{ui})^2 + \frac{\alpha}{2}\|U_{u\cdot}\|^2 + \frac{\alpha}{2}\|V_{i\cdot}\|^2 + \frac{\alpha}{2}\|b_u\|^2 + \frac{\alpha}{2}\|b_i\|^2 + \delta_p \frac{\alpha}{2} \sum_{j \in \mathcal{P}_u} \|C_{j\cdot}^{p}\|_F^2 + \delta_n \frac{\alpha}{2} \sum_{j \in \mathcal{N}_u} \|C_{j\cdot}^{n}\|_F^2 + \delta_{\mathcal{O}} \frac{\alpha}{2} \sum_{i' \in \mathcal{I}_u \backslash \{i\}} \|C_{i'\cdot}^{o}\|_F^2 + \delta_{\mathcal{G}} \frac{\alpha}{2} \sum_{g \in \mathbb{G}} \sum_{i' \in \mathcal{I}_u^g \backslash \{i\}} \|C_{i'\cdot}^{g}\|_F^2$, and $\tilde{f}_{ui} = \frac{1}{2}(\tilde{r}_{ui} - \hat{\tilde{r}}_{ui})^2 + \frac{\alpha}{2}\|W_{u\cdot}\|^2 + \frac{\alpha}{2}\|V_{i\cdot}\|^2$.

For solving the optimization problem in (7), we randomly take a record from $\mathcal{R} \cup \tilde{\mathcal{R}}$, and update the corresponding model parameters in Θ, i.e., μ, b_u, b_i, $U_{u\cdot}$, $V_{i\cdot}$, $C_{j\cdot}^{p}$, $C_{j\cdot}^{n}$, $C_{i'\cdot}^{o}$ and $C_{i'\cdot}^{g}$.

TABLE I
SOME NOTATIONS (INCLUDING THOSE OF THE DATA, MODELS AND LEARNING ALGORITHMS) AND EXPLANATIONS USED IN THE PAPER.

n	number of users
m	number of items
$u, u' \in \{1, 2, \ldots, n\}$	user ID
$i, j \in \{1, 2, \ldots, m\}$	item ID
$\mathbb{G} = \{0.5, 1, \ldots, 5\}$	grade score range
$\mathbb{B} = \{\text{like}, \text{dislike}\}$	binary rating range
$r_{ui} \in \mathbb{G}$	grade score of user u to item i
$\tilde{r}_{ui} \in \mathbb{B}$	binary rating of user u to item i
$\mathcal{R} = \{(u, i, r_{ui})\}$	grade score records (training data)
$\tilde{\mathcal{R}} = \{(u, i, \tilde{r}_{ui})\}$	binary rating records (training data)
$p = \lvert\mathcal{R}\rvert$	number of grade scores (training data)
$\tilde{p} = \lvert\tilde{\mathcal{R}}\rvert$	number of binary ratings (training data)
$\mathcal{I}_u^g, g \in \mathbb{G}$	items rated by user u with score g (training data)
$\mathcal{P}_u$	items liked (w/ positive feedback) by user u (training data)
$\mathcal{N}_u$	items disliked (w/ negative feedback) by user u (training data)
$\mathcal{T}_E = \{(u, i, r_{ui})\}$	grade score records in test data
$\mu \in \mathbb{R}$	global average rating value
$b_u \in \mathbb{R}$	user bias
$b_i \in \mathbb{R}$	item bias
$d \in \mathbb{R}$	number of latent dimensions
$U_{u\cdot}, W_{u\cdot} \in \mathbb{R}^{1\times d}$	user-specific latent feature vector
$\mathbf{U}, \mathbf{W} \in \mathbb{R}^{n\times d}$	user-specific latent feature matrix
$V_{i\cdot}, C^p_{j\cdot}, C^n_{j\cdot}, C^o_{i'\cdot}, C^g_{i'\cdot} \in \mathbb{R}^{1\times d}$	item-specific latent feature vector
$\mathbf{V}, \mathbf{C}^p, \mathbf{C}^n, \mathbf{C}^o, \mathbf{C}^g \in \mathbb{R}^{m\times d}$	item-specific latent feature matrix
$\hat{r}_{ui}$	predicted grade score of user u to item i
$\hat{\tilde{r}}_{ui}$	predicted binary rating of user u to item i
γ	learning rate
ρ	interaction weight between grade scores and binary ratings
α	tradeoff parameter on the regularization terms
$\delta_{\mathcal{O}}, \delta_{\mathcal{G}}, \delta_p, \delta_n \in \{0, 1\}$	indicator variable for positive and negative feedback
w_p, w_n	weight on positive and negative feedback
T	iteration number in the algorithm

if a grade score record $(u, i, r_{ui}) \in \mathcal{R}$ is sampled (i.e., $y_{ui} = 1$), and $W_{u\cdot}$ and $V_{i\cdot}$ if a binary rating record $(u, i, \tilde{r}_{ui}) \in \tilde{\mathcal{R}}$ is sampled (i.e., $\tilde{y}_{ui} = 1$). In order to further introduce rich interactions between the two factorization subtasks, we follow the technique in iTCF [13], which will be reflected in the gradients in the sequel.

Specifically, for a randomly picked up record, we have the gradients of the model parameters w.r.t. f_{ui} or $\tilde{f}_{ui}$ as follows,

$$
\begin{aligned}
\nabla\mu &= -e_{ui} \\
\nabla b_u &= -e_{ui} + \alpha b_u \\
\nabla b_i &= -e_{ui} + \alpha b_i \\
\nabla U_{u\cdot} &= -e_{ui} V_{i\cdot} + \alpha U_{u\cdot} \\
\nabla V_{i\cdot} &= -e_{ui}(\rho U_{u\cdot} + (1-\rho) W_{u\cdot} + \bar{C}^p_{u\cdot} + \bar{C}^n_{u\cdot} + \bar{C}^{\mathcal{G}}_{u\cdot}) + \alpha V_{i\cdot} \\
\nabla C^p_{j\cdot} &= \delta_p(-e_{ui} w_p \frac{1}{\sqrt{|\mathcal{P}_u|}} V_{i\cdot} + \alpha C^p_{j\cdot}), j \in \mathcal{P}_u \\
\nabla C^n_{j\cdot} &= \delta_n(-e_{ui} w_n \frac{1}{\sqrt{|\mathcal{N}_u|}} V_{i\cdot} + \alpha C^n_{j\cdot}), j \in \mathcal{N}_u \\
\nabla C^o_{i'\cdot} &= \delta_{\mathcal{O}}(-e_{ui} \frac{1}{\sqrt{|\mathcal{I}_u \backslash \{i\}|}} V_{i\cdot} + \alpha C_{i'\cdot}), i' \in \mathcal{I}_u \backslash \{i\} \\
\nabla C^g_{i'\cdot} &= \delta_{\mathcal{G}}(-e_{ui} \frac{1}{\sqrt{|\mathcal{I}^g_u \backslash \{i\}|}} V_{i\cdot} + \alpha C^g_{i'\cdot}), i' \in \mathcal{I}^g_u \backslash \{i\}, g \in \mathbb{G} \\
\nabla W_{u\cdot} &= \lambda(-\tilde{e}_{ui} V_{i\cdot} + \alpha W_{u\cdot}) \\
\nabla V_{i\cdot} &= \lambda(-\tilde{e}_{ui}(\rho W_{u\cdot} + (1-\rho) U_{u\cdot}) + \alpha V_{i\cdot})
\end{aligned}
$$

where $e_{ui} = (r_{ui} - \hat{r}_{ui})$ is the error w.r.t. the target grade score, $\tilde{e}_{ui} = (\tilde{r}_{ui} - \hat{\tilde{r}}_{ui})$ is the error w.r.t. the auxiliary binary rating, and $\rho U_{u\cdot} + (1-\rho) W_{u\cdot}$ is used to introduce rich interactions [13] between the user-specific latent features $U_{u\cdot}$ and $W_{u\cdot}$. Notice that we follow TMF [8] and convert each "like" to $\tilde{r}_{ui} = 5$ and each "dislike" to $\tilde{r}_{ui} = 1$. In our factorization-based transfer learning framework, the implicit preference context is extracted from both grade scores and binary ratings, which are both leveraged to improve the task of predicting grade scores. On the one hand, the explicit grade scores themselves are an obvious symbol of classification which is easily neglected but important to capture users' detailed information. On the other hand, the binary ratings indicate users' affinity or rejection directly, which helps to model users' intuitive feelings. More importantly, by combining them in one single framework, knowledge can be transferred between each other so that contributes to more accurate modeling of users' preferences.

D. Discussions

Our transfer learning solution is very generic, which contains several state-of-the-art factorization-based recommendation methods as special cases, including RSVD [10], CMF [14], iTCF [13], TMF [8], SVD++ [10] and MF-MPC [11]. We explicitly describe the relationship of these models in Table II. From Table II, we can see that our PAT contains several pluggable components such as the components for the auxiliary binary ratings, the positive and negative preference context, the multiclass preference context, and the interaction between two types of feedback, which shows that our solution is very generic and flexible. This also gives us a new perspective of analyzing the performance of each specific method, which is included in our empirical studies.

TABLE II
RELATIONSHIPS BETWEEN OUR PREFERENCE-AWARE TRANSFER (PAT) AND OTHER FACTORIZATION-BASED METHODS IN THE PERSPECTIVE OF ITS PROJECTION TO THE GRAPHICAL MODEL OF OUR PAT SHOWN IN FIG. 1.

Algorithm	Edges
RSVD [10]	$\{e_1, e_2\}$
CMF [14]	$\{e_1, e_2, e_3, e_4\}$
iTCF [13]	$\{e_1, e_2, e_3, e_4, e_5\}$
TMF [8]	$\{e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$
SVD++ [10], MF-MPC [11]	$\{e_1, e_2, e_8\}$
PAT-OPC, PAT (**proposed**)	$\{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_8\}$

E. The Learning Algorithm

We adopt the commonly used stochastic gradient descent (SGD) based algorithm to solve the optimization problem, and formally describe the algorithm in Algorithm 1. We can see that our PAT algorithm mainly contains two loops, where the core steps are as follows: (i) sampling a record from a union of two sets $\mathcal{R} \cup \tilde{\mathcal{R}}$, (ii) calculating the gradients of the corresponding model parameters, and (iii) updating the model parameters. The time complexity of our PAT are comparable to that of TMF [8] and MF-MPC [11].

Algorithm 1 The algorithm of preference-aware transfer (PAT).

```
for t = 1, ..., T do
    for iter = 1, ..., |R ∪ R̃| do
        Randomly pick up a record (u, i, r_ui) or (u, i, r̃_ui)
        from R ∪ R̃.
        Calculate the gradients w.r.t. f_ui or f̃_ui accordingly.
        Update the corresponding model parameters.
    end for
    Decrease the learning rate via γ ← γ × 0.9.
end for
```

IV. EXPERIMENTAL RESULTS

In this section, we conduct empirical studies in order to study the usefulness of the binary rating data and the effectiveness of exploiting the preference context in terms of the prediction accuracy of the grade scores.

A. Datasets and Evaluation Metrics

In order to compare with the closely related works directly, we adopt two public datasets used in a previous study about modeling heterogeneous feedback [8], i.e, Movielens 10M (denoted as ML10M) and Flixter. The ML10M dataset contains $10,000,054$ grade scores from $71,567$ users to $10,681$ items. The Flixter dataset contains $8,196,075$ grade scores from $147,612$ users to $48,794$ items. For simulating the problem setting with heterogeneous feedback, we process each dataset as follows: (i) we randomly split the data into five parts with similar size, and (ii) we then take two parts as training data with grade scores, take another two parts as binary ratings by transforming grade scores larger than or equal to four to "like" and grade scores smaller than four to "dislike", and take the remaining one part as the test data with grade scores. We repeat this process for five times and obtain five copies of grade score records, binary ratings and test data. The results are averaged over those five copies of data.

To evaluate the performance of our PAT and other methods, we use two commonly used regression-oriented evaluation metrics, which are mean absolute error (MAE) and root mean square error (RMSE).

B. Baselines and Parameter Configuration

In order to study the effectiveness of our generic transfer learning framework, we conduct the experiments with the following baseline methods:

- RSVD [10] is a basic matrix factorization method without modeling preference context and auxiliary binary ratings, which is a special case of our PAT with edges $\{e_1, e_2\}$;
- MF-MPC [11] is a recent advanced matrix factorization method exploiting multiclass preference context beneath the grade scores, which is a special case of our PAT with edges $\{e_1, e_2, e_8\}$; and
- TMF [8] is a recent factorization-based transfer learning method incorporating the auxiliary binary ratings, which is a special case of our PAT with edges $\{e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$.

Notice that we do not include some other algorithms for the studied problem such as collective matrix factorization (CMF) [14] with edges $\{e_1, e_2, e_3, e_4\}$, and interaction-rich transfer by collective facrtorization (iTCF) [13] with edges $\{e_1, e_2, e_3, e_4, e_5\}$ because they usually perform worse than TMF [8].

As for the parameter configurations in our PAT, we adhere to the same rules used in TMF [8]. Specially, we fix the number of latent dimensions $d = 20$ on ML10M and $d = 10$ on Flixter, respectively, the iteration number $T = 50$, the learning rate $\gamma = 0.01$, the interaction weight $\rho = 0.5$, the weight on the auxiliary binary ratings $\lambda = 1$, the tradeoff parameter on the regularization terms $\alpha = 0.01$, and the weight on positive and negative feedback $w_p = 2$ and $w_n = 1$.

For fair comparison between our PAT and other factorization-based methods, we use the same code framework of our PAT in Java programming language, and obtain the

TABLE III

RECOMMENDATION PERFORMANCE OF OUR PREFERENCE-AWARE TRANSFER (PAT) AND OTHER FACTORIZATION-BASED METHODS ON ML10M AND FLIXTER, WHERE THE RESULTS OF RSVD [10] AND TMF [8] ARE COPIED FROM [8]. NOTICE THAT WE FOLLOW THE PARAMETER SETTING IN TMF [8] AND FIX $\alpha = 0.01$ AND $T = 50$ FOR ALL THE METHODS, AND $w_p = 2$ AND $w_n = 1$ FOR TMF AND OUR PAT. WE ALSO INCLUDE THE CONFIGURATIONS IN OUR GENERIC PAT FRAMEWORK FOR COMPARATIVE STUDY AND REPRODUCIBILITY.

Data	Algorithm	MAE	RMSE	Configurations
ML10M	RSVD	0.6438± 0.0011	0.8364± 0.0012	$\delta_p = \delta_n = 0, \delta_{\mathcal{G}} = 0, \lambda = 0, \rho = 1$
	MF-MPC	0.6162± 0.0006	0.8063± 0.0007	$\delta_p = \delta_n = 0, \delta_{\mathcal{G}} = 1, \lambda = 0, \rho = 1$
	TMF	0.6124± 0.0007	0.8005± 0.0008	$\delta_p = \delta_n = 1, \delta_{\mathcal{G}} = 0, \lambda = 1, \rho = 0.5$
	PAT	**0.6107**± 0.0003	**0.7989**± 0.0008	$\delta_p = \delta_n = 1, \delta_{\mathcal{G}} = 1, \lambda = 1, \rho = 0.5$
Flixter	RSVD	0.6561± 0.0007	0.8814± 0.0010	$\delta_p = \delta_n = 0, \delta_{\mathcal{G}} = 0, \lambda = 0, \rho = 1$
	MF-MPC	0.6383± 0.0004	0.8644± 0.0005	$\delta_p = \delta_n = 0, \delta_{\mathcal{G}} = 1, \lambda = 0, \rho = 1$
	TMF	0.6348± 0.0007	0.8615± 0.0012	$\delta_p = \delta_n = 1, \delta_{\mathcal{G}} = 0, \lambda = 1, \rho = 0.5$
	PAT	**0.6332**± 0.0006	**0.8572**± 0.0010	$\delta_p = \delta_n = 1, \delta_{\mathcal{G}} = 1, \lambda = 1, \rho = 0.5$

baseline methods via specific configurations, which are also included in our experimental results in Table III.

C. Results

The recommendation performance of our PAT and other factorization-based methods are shown in Table III, from which we can have the following observations:

- our PAT performs significantly better than all the baseline methods across the two datasets, which shows the effectiveness of our transfer learning solution in modeling users' heterogeneous feedback and preference context; and
- compared with RSVD and MF-MPC, TMF and our PAT with both target grade scores and auxiliary binary ratings perform better, which showcases the usefulness of the binary ratings.

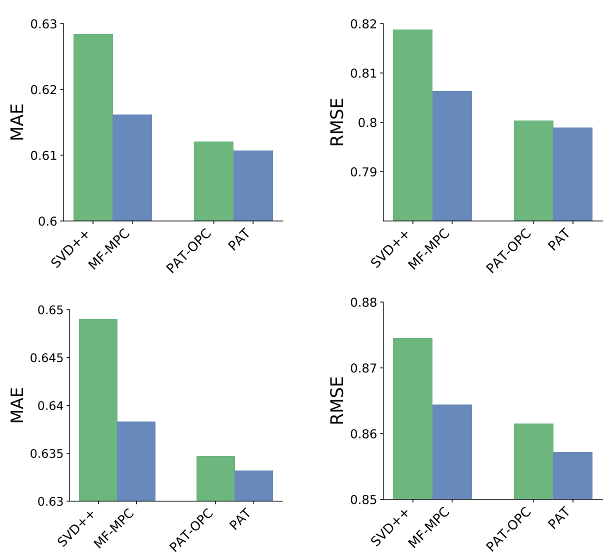

Fig. 2. Recommendation performance of factorization methods with one-class preference context (OPC) and multiclass preference context (MPC), i.e., MF with OPC (SVD++ [10]), MF with MPC (MF-MPC [11]), reduced version of our PAT with OPC (i.e., PAT-OPC) and our PAT with MPC (PAT) on ML10M (top) and Flixter (bottom), respectively.

We also study the effectiveness of oneclass preference context (OPC) and multiclass preference context (MPC) in factorization-based methods in depth, which are shown in Fig. 2. In particular, we include basic matrix factorization with OPC and MPC denoted as SVD++ [10] and MF-MPC [11], respectively, and a reduced version of our PAT with OPC (i.e., PAT-OPC) and our PAT. Technically, PAT-OPC can easily be derived and implemented by removing the grade scores when generating the multiclass preference context, i.e., $g = 1$ for all values in $\{0.5, 1, \ldots, 5\}$. Notice that SVD++ and MF-MPC only exploit the grade score records $\mathcal{R}$, while our PAT-OPC and PAT make use of both the grade scores and the binary ratings. From the results in Fig. 2, we can have the following observations:

- the overall performance ordering is SVD++ $<$ MF-MPC $<$ PAT-OPC $<$ PAT, which clearly showcases the effectiveness of our preference-aware transfer learning solution in modeling users' heterogeneous feedback;
- for the two methods with OPC, i.e., SVD++ and our PAT-OPC, and the two methods with MPC, i.e., MF-MPC and our PAT, we can see that integrating the binary rating records always brings performance improvement; and
- for the two methods modeling grade scores only, i.e., SVD++ and MF-MPC, and the two methods modeling both grade scores and binary ratings, i.e., PAT-OPC and PAT, we find that a method with MPC is always better than that with OPC, which shows the effectiveness of leveraging the fine-granularity preference context.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we propose a novel and generic transfer learning framework, i.e., preference-aware transfer (PAT), for recommendation with heterogeneous feedback. In particular, we take the grade scores as the target data and the likes/dislikes as the auxiliary data in a transfer learning view, and exploit the implicit preference beneath the target data and the auxiliary data as the preference context, in order to build a more accurate and generic recommendation model. Technically, we find that several recent algorithms can be projected to be parts of our generic solution PAT as special cases. Empirically, we obtain very promising results on two large and public datasets in comparison with several state-of-the-art methods. More importantly, we observe that the empirical results are consistent with that of the technical framework with different subsets of components, i.e., more components leading to better performance.

For future works, we are interested in further generalizing our generic factorization framework with deep federated learning [15], [16] and ranking-oriented recommendation [17], [18].

Acknowledgment

We thank the support of National Natural Science Foundation of China Nos. 61872249, 61836005 and 61672358. Weike Pan and Zhong Ming are corresponding authors for this work.

References

[1] M. Zihayat, A. Ayanso, X. Zhao, H. Davoudi, and A. An, "A utility-based news recommendation system," *Decision Support Systems*, vol. 117, pp. 14–27, 2019.

[2] B. De Carolis, M. de Gemmis, P. Lops, and G. Palestra, "Recognizing users feedback from non-verbal communicative acts in conversational recommender systems," *Pattern Recognition Letters*, vol. 99, pp. 87–95, 2017.

[3] D. Horowitz, D. Contreras, and M. Salamó, "Eventaware: A mobile recommender system for events," *Pattern Recognition Letters*, vol. 105, pp. 121–134, 2018.

[4] Y. Shi, M. Larson, and A. Hanjalic, "Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges," *ACM Computing Surveys (CSUR)*, vol. 47, no. 1, p. 3, 2014.

[5] Y. Huang and J. T. Kwok, "Collaborative filtering via co-factorization of individuals and groups," in *Proceeding of the 2015 International Joint Conference on Neural Networks*, 2015, pp. 1–8.

[6] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, no. 8, pp. 30–37, 2009.

[7] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," in *Annual Conference on Neural Information Processing Systems*, 2008, pp. 1257–1264.

[8] W. Pan, S. Xia, Z. Liu, X. Peng, and Z. Ming, "Mixed factorization for collaborative recommendation with heterogeneous explicit feedbacks," *Information Sciences*, vol. 332, pp. 84–93, 2016.

[9] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[10] Y. Koren, "Factorization meets the neighborhood: A multifaceted collaborative filtering model," in *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008, pp. 426–434.

[11] W. Pan and Z. Ming, "Collaborative recommendation with multiclass preference context," *IEEE Intelligent Systems*, vol. 32, no. 2, pp. 45–51, 2017.

[12] W. Pan and Q. Yang, "Transfer learning in heterogeneous collaborative filtering domains," *Artificial Intelligence*, vol. 197, pp. 39–55, 2013.

[13] W. Pan and Z. Ming, "Interaction-rich transfer learning for collaborative filtering with heterogeneous user feedbacks," *IEEE Intelligent Systems,*, vol. 29, no. 6, pp. 48–54, 2014.

[14] A. P. Singh and G. J. Gordon, "Relational learning via collective matrix factorization," in *Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008, pp. 650–658.

[15] F. Xue, X. He, X. Wang, J. Xu, K. Liu, and R. Hong, "Deep item-based collaborative filtering for top-n recommendation," *ACM Transactions on Information Systems*, vol. 37, no. 3, pp. 33:1–33:25, 2019.

[16] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 12:1–12:19, 2019.

[17] L. Wu, C. Hsieh, and J. Sharpnack, "SQL-Rank: A listwise approach to collaborative ranking," in *Proceedings of the 35th International Conference on Machine Learning*, ser. ICML '18, 2018, pp. 5311–5320.

[18] C. Pei, Y. Zhang, Y. Zhang, F. Sun, X. Lin, H. Sun, J. Wu, P. Jiang, J. Ge, W. Ou, and D. Pei, "Personalized re-ranking for recommendation," in *Proceedings of the 13th ACM Conference on Recommender Systems*, 2019, pp. 3–11.