Expert Systems With Applications 74 (2017) 11–18

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Preference dynamics with multimodal user-item interactions in social media recommendation

D. Rafailidis*, P. Kefalas, Y. Manolopoulos

Department of Informatics, Aristotle University, Thessaloniki 54124, Greece

ARTICLE INFO

Article history:
Received 22 September 2016
Revised 4 January 2017
Accepted 4 January 2017
Available online 5 January 2017

Keywords:
Recommender systems
Multimodal information
Preference dynamics
Collective matrix factorization

ABSTRACT

Recommender systems elicit the interests and preferences of individuals and make recommendations accordingly, a main challenge for expert and intelligent systems. An essential problem in recommender systems is to learn users' preference dynamics, that is, the constant evolution of the explicit or the implicit information, which is diversified throughout time according to the user actions. Also, in real settings data sparsity degrades the recommendation accuracy. Hence, state-of-the-art methods exploit multimodal information of users-item interactions to reduce sparsity, but they ignore preference dynamics and do not capture users' most recent preferences. In this article, we present a *Temporal Collective Matrix Factorization* (*TCMF*) model, making the following contributions: (i) we capture preference dynamics through a joint decomposition model that extracts the user temporal patterns, and (ii) co-factorize the temporal patterns with multimodal user-item interactions by minimizing a joint objective function to generate the recommendations. We evaluate the performance of *TCMF* in terms of accuracy and root mean square error, and show that the proposed model significantly outperforms state-of-the-art strategies.

1. Introduction

Collaborative Filtering (CF) plays a vital role in recommender systems, where users with the same activity, for example, rating or tagging behavior, tend to get similar recommendations (Sarwar, Karypis, Konstan, & Riedl, 2001; Tso, Marinho, & Schmidt-Thieme, 2008). In real-world recommender systems, users tend to change their behavior temporally due to the evolving preference propagation (Koren, 2009). To face this problem, many methods capturing the users' preference dynamics have been presented, but they suffer from the data sparsity, having only a few user-item interactions to base the recommendations. Other methods face the data sparsity by exploiting the multimodal information of user-item interactions. However they do not capture users' preference dynamics which are crucial as preferences evolve over time. To the best of our knowledge, there is no model in the literature that simultaneously captures users' preference dynamics and multimodal information about user-item interactions. Here, we present a *Temporal Collective Matrix Factorization* model, namely *TCMF*, that combines both users' preference dynamics and multimodal information of users-item interactions. Next, we briefly discuss how users' preference dynamics and the exploitation of multimodal user-item interactions impact the recommendation accuracy.

1.1. Users' preference dynamics

Users' preference dynamics have a serious impact on recommender systems because they limit the recommendation accuracy as the recent users' preferences are ignored. It is also widely known in literature as *temporal dynamics* (Zhang, Wang, Yu, Sun, & Lim, 2014a). Researchers pointed out that the reasons why users change their taste are (Koren, 2009; Lathia, Hailes, Capra, & Amatriain, 2010; Xiong, Chen, Huang, Schneider, & Carbonell, 2010; Yin, Gupta, Weninger, & Han, 2010): (i) *New items exploration:* curiosity leads users to explore new items contrary to their ordinary choices; (ii) *User experience:* if a user has a pleasant experience in the past, then probably s/he will choose the same or a similar interaction in the future; (iii) *Popularity:* users interact with a bias based on popularity irrespective to their history record; (iv) *Social influence:* friends' opinion is essential while making decisions, where users tend to examine their friends' evaluation and follow their lead; (v) *User expertise:* users become more and more familiar with items they interact and thus gain meaningful experience, i.e., they become skeptical and ungenerous while rating similar items to the ones rated in the past. For all these reasons, users' preference dynamics should be taken into consideration when generating the recommendations.

* Corresponding author.
E-mail addresses: draf@csd.auth.gr, drafail@iti.gr (D. Rafailidis), kefalasp@csd.auth.gr (P. Kefalas), manolopo@csd.auth.gr (Y. Manolopoulos).

http://dx.doi.org/10.1016/j.eswa.2017.01.005
0957-4174/

1.2. Multimodal user-item interactions

Multimodal information of user-item interactions concerns data that explicitly or implicitly refer to the user actions (Saveski & Mantrach, 2014). In contrast to implicit data that are computed without prior knowledge about users past history, such as statistical scores, views, user influence and other, explicit data are provided by users themselves straightforwardly, such as ratings, comments, and tags. The main difference is that implicit information represents knowledge about user statistical profile, whereas explicit information enhances and personalizes our knowledge about the user actual preferences. CF relies on user past history to make predictions, thus it is difficult to provide a recommendation for a user with a short history record. This problem is known as the cold start problem (Herlocker, Konstan, Terveen, & Riedl, 2004), which limits the accuracy of systems based on the CF method by magnifying the data sparsity problem. Several approaches have been proposed in the literature to deal with the cold start and the data sparsity problems, such as: (i) *hybrid models* that combine CF and content-based methods (Jialu, Chi, Gao, & Jiawei, 2013; Lee & Seung, 1999); (ii) *new similarity measures* (Ahn, 2008; Bobadilla, Ortega, Hernando, & Bernal, 2012); and (iii) *co-factorizing*[1] *or coupling information from auxiliary sources* (Narita, Hayashi, Tomioka, & Kashima, 2011; Rafailidis & Nanopoulos, 2014). According to co-factorizing techniques, the multimodal information of user-item interactions, provided either implicitly or explicitly by the user themselves, improves the recommendation accuracy by enriching the knowledge about user preferences.

1.3. Contribution and outline

Summarizing, state-of-the-art strategies either do not capture users' preference dynamics or exploit multimodal information of user-item interactions, but not both of them simultaneously. For instance, methods which co-factorize or couple multimodal information (Li, Liu, Zhu, & Lu, 2010; Narita et al., 2011) have the advantage that they face the data sparsity problem, but have the disadvantage that they do not capture the temporal dimension. Therefore, they miss users' preference evolution, a crucial factor in recommender systems, as users tent to change their behavior and their taste over time (Cho, Myers, & Leskovec, 2011; Koren, 2009). On the other hand, methods that capture the temporal dimension, such as the studies reported in Koren (2009); Vaca, Mantrach, Jaimes, and Saerens (2014); Zhang et al. (2014a) allow to capture users' preference dynamics, but have the disadvantage that they do not exploit multimodal information resulting in limited recommendations accuracy by not handling the data sparsity problem.

In this article, we introduce the *TCMF* model. Our main contribution is summarized as follows: **(C1)** *TCMF* captures users' preference dynamics through a joint decomposition model that extracts the user temporal pattern. Our model minimizes a loss function by using multiplicative updating rules for both past and future time periods. **(C2)** Our model co-factorizes the user temporal patterns with multimodal user-item interactions by defining a joint objective function to generate personalized recommendations. **(C3)** We evaluate the performance of our method against state-of-the-art algorithms with respect to the recommendation accuracy and root mean squared error, and we show that the proposed *TCMF* model is superior over the state-of-the art algorithms on benchmark data sets that span from 13 to 15 years that also contain multimodal information of user-item interactions.

[1] The term co-factorize refers to the joint factorization of multiple matrices, each one representing a relation between two entity types, for example, multimodal user-item interactions (Bouchard, Yin, & Guo, 2013; Huang, Xiang, & Pan, 2012; Singh & Gordon, 2008).

The remainder of this article is organized as follows, Section 2 summarizes the related studies, whereas Section 3 introduces the mathematical formulation of our problem. Section 4 illustrates our method structural parts in detail and introduces our algorithm. Experimental results are given in Section 5. Finally, Section 6 concludes this paper.

2. Related work

In many recent studies efforts have been made towards: (i) capturing users' preference dynamics (Koren, 2009; Vaca et al., 2014; Zhang et al., 2014a); and (ii) exploiting multimodal information to provide recommendations (Caicedo, BenAbdallah, González, & Nasraoui, 2012; Da Costa, Domingues, Rezende, & Manzato, 2014; Saveski & Mantrach, 2014).

2.1. Capturing users' preference dynamics

Koren (2009) introduces the *timeSVD*++ algorithm, which captures the lasting and the transient factors by modeling the users' preference dynamics through the entire time period. The goal is to distill longer term preferences from the noisy patterns using a matrix factorization model. He shows that in an item-item neighborhood model, the essential relations among the items can be extracted by learning how ratings evolve. Similarly, Zhang et al. (2014a) present two models which capture the users' preference dynamics concerning users' preference evolution, namely the Temporal Matrix Factorization (*TMF*) and the Bayesian Temporal Matrix Factorization (*BTMF*) methods. *TMF* maps users' preferences on items into a joint latent factor with a transition matrix, which captures the users' preference dynamics between two time periods. Thus, they sample the rating distribution from the users and items, to update the transition matrix for both past and future time periods. *BTMF* extends *TMF* by introducing priors for the hyper parameters to increase the accuracy and deals with the complexity of *TMF*. Vaca et al. (2014) introduce a collective matrix factorization model, which tracks emerging, evolving and fading latent factors according to the past history. This model learns jointly the preferences evolution and their time dependencies using a mapping matrix to capture the users' preference dynamics between two time periods. This matrix depicts the diversification between two time periods and captures the evolution of the distributions for all involved entities. However, this work focuses on topic discovery and monitoring news and not on generating recommendations.

Yin, Cui, Chen, Hu, and Zhou (2015) point out that users' rating behavior is influenced by two factors: users' preferences and social preferences. Also, they show that both factors have different degree of influence on user rating behavior. To identify the users' rating behavior affected by both factors they present a Temporal Context-Aware Recommender System (*TCARS*), which jointly exploits both the users' preferences and social preferences at different time periods. Similarly, Liu, He, and Zhao (2013) present a Social Temporal Collaborative Ranking ($ST-CoR$) model to provide item and user recommendations. $ST-CoR$ aggregates heterogeneous, explicit and implicit information of users' feedback such as ratings, comments, favorites and collections. In addition, they extend this model to capture behaviors that are constant versus those which are not. Due to limitations of the user feedback, they incorporate social-based similarity to the model to overcome the sparsity problem. However, both *TCARS* and $ST-CoR$ incorporate users' social relations with information from heterogeneous networks, which is out of this paper's scope and is left for future work.

Table 1
Notation.

Symbol	Description
m, n, d	number of items, users and latent factors
$X^{(t)}, X^{(t-1)} \in \mathbb{R}^{m\times n}$	target matrices in the present and past time periods t and $t-1$
$A^{(t)}, A^{(t-1)} \in \mathbb{R}^{m\times n}$	auxiliary matrices in t and $t-1$
$\hat{X}^{(t)}, \hat{X}^{(t-1)} \in \mathbb{R}^{m\times n}$	factorized target matrices of $X^{(t)}$ and $X^{(t-1)}$
$\hat{A}^{(t)}, \hat{A}^{(t-1)} \in \mathbb{R}^{m\times n}$	factorized auxiliary matrices of $A^{(t)}$ and $A^{(t-1)}$
$Z_x^{(t)} \in \mathbb{R}^{m\times d}, V_x^{(t)} \in \mathbb{R}^{d\times n}$	item and user factor matrices of the target matrix $X^{(t)}$
$Z_a^{(t)} \in \mathbb{R}^{m\times d}, V_a^{(t)} \in \mathbb{R}^{d\times n}$	item and user factor matrices of the auxiliary matrix $A^{(t)}$
$T_x^{(t)} \in \mathbb{R}^{d\times d}$	transition matrix between user factor matrices $V_x^{(t-1)}$ and $V_x^{(t)}$
$T_a^{(t)} \in \mathbb{R}^{d\times d}$	transition matrix between user factor matrices $V_a^{(t-1)}$ and $V_a^{(t)}$
λ	temporal regularization parameter
β	l_1-norm regularization parameter for the factor matrices
γ	l_1-norm regularization parameter for the transition matrices

2.2. Exploiting multimodal information

Unlikely to the above temporal-based methods, several methods incorporate side information to improve the recommendation accuracy by handling the data sparsity problem, such as: (i) *multimodal information of user-item interactions*, such as scores, views, users' influence, ratings, comments and other (Caicedo et al., 2012; Da Costa et al., 2014; Da Costa & Manzato, 2014; Saveski & Mantrach, 2014); (ii) *side information about users' attributes* such as demographics, gender, features, text and other attributes (Fang & Si, 2011; Narita et al., 2011; Porteous, Asuncion, & Welling, 2010; Zhou, Shan, Banerjee, & Sapiro, 2012); and (iii) *content information about item features* (Kim, Kim, Lee, & Lee, 2014; Liu et al., 2013; Yin et al., 2015; Zhang, Li, Chen, Zhang, & Wang, 2014b). In the context of our research we focus on multimodal information of user-item interactions and not on side information, where users must provide this information themselves, and not on content information that requires a heavy computational analysis of data, which may not always be feasible for millions of items.

Saveski and Mantrach (2014) propose a Local Collective Embeddings (*LCE*) model to tackle the sparsity problem. *LCE* learns items properties and user preferences similarity simultaneously. This model follows an item-based approach to provide recommendations, accounting for the fact that a new item with its description predicts the users who are likely to be interested in this item. However, *LCE* fails to capture the evolution of users' preferences. In the studies reported in Rafailidis and Nanopoulos (2014, 2015), the models capture the rate with which the past preferences of each user have been shifted temporally at the current time period, and then the importance of user preferences is weighted accordingly. These models examine the time dimension and side information through similarity measures but concern repeat consumption recommendations and not novel ones, as this study does.

3. Problem formulation

Table 1 presents the main notation used in the sequel of the article.

Let t be the present time period, on which we have to generate the recommendations. Also, let $t-1$ be the past time period, that is, the accumulated history. The choice of the time periods, for example days, weeks or months depends on how often the recommendations are generated in the social media platform. For example, given a total period of ten months, the present period t is the tenth month and the past time periods are the last nine months. Given m items and u users, we assume that users express their preferences over items in different types of interactions, such as explicit interaction e.g., ratings, tag assignments or implicit one e.g., number of views, clicks, comments and so on. This means that we have a target user-item interaction type, where the recommendations are generated, and different auxiliary interaction types. To simplify the presentation, we consider two types of user-item interactions, stored in a target matrix $X^{(t)} \in \mathbb{R}^{m\times n}$ and an auxiliary matrix $A^{(t)} \in \mathbb{R}^{m\times n}$. The entries $[X^{(t)}]_{ij}$ and $[A^{(t)}]_{ij}$ denote the target and the auxiliary interactions of user j over item i in the present time period t, respectively. To account for the accumulated history of the target and auxiliary interactions, we also consider the interaction matrices $X^{(t-1)}$ and $A^{(t-1)}$ in the past period $t-1$.

Following the Non-Negative Matrix Factorization (NMF) technique, introduced in Lee and Seung (2000), for a single type of user-item interactions the recommendations are generated in the present time period t by factorizing the target matrix $X^{(t)}$. This is achieved by decomposing $X^{(t)}$ as follows:

$$X^{(t)} \approx Z_x^{(t)} V_x^{(t)} \tag{1}$$

subject to $Z_x^{(t)}, V_x^{(t)} \geq 0$, where $Z_x^{(t)} \in \mathbb{R}^{m\times d}$ and $V_x^{(t)} \in \mathbb{R}^{d\times n}$ are the factor matrices of items and users, respectively, and d is the number of latent factors. The ith row of $Z_x^{(t)}$ and the jth column of $V_x^{(t)}$ express the d-dimensional latent vectors of item i and user j, accordingly. The product $Z_x^{(t)} V_x^{(t)}$ results in a factorized matrix $\hat{X}^{(t)}$, that is, the low rank d approximation of $X^{(t)}$, with $X^{(t)} \approx \hat{X}^{(t)}$.

In our setting, we have to calculate the factorized matrix $\hat{X}^{(t)} = Z_x^{(t)} V_x^{(t)}$, by also considering the users' preference dynamics, that is, the transition of user preferences between the past and the present time periods $t-1$ and t, as well as the auxiliary user-item interactions in these periods. Our problem is formally defined as follows:

Definition 1 (Problem)**.** "Given (i) the target interaction matrices $X^{(t)}$ and $X^{(t-1)}$ and (ii) the auxiliary matrices $A^{(t)}$ and $A^{(t-1)}$ in the time periods t and $t-1$, the goal of the proposed approach is to compute the factorized matrix $\hat{X}^{(t)}$, by capturing both the preference dynamics and the multimodal user-item interactions".

4. Proposed approach

In Section 4.1, first we formulate the joint objective function as a minimization problem to capture both the preference dynamics and multimodal user-item interaction, and then in Section 4.2 we detail the optimization algorithm to solve the joint minimization problem, to generate the final recommendations in the factorized matrix $\hat{X}^{(t)}$.

4.1. Objective function

The NMFs of matrices $X^{(t)}$ and $X^{(t-1)}$ for the periods t and $t-1$ correspond to the following minimization problems:

$$\min_{Z_x^{(t)}, V_x^{(t)}} ||X^{(t)} - Z_x^{(t)} V_x^{(t)}||_F^2 \tag{2}$$

$$\min_{Z_x^{(t-1)}, V_x^{(t-1)}} ||X^{(t-1)} - Z_x^{(t-1)} V_x^{(t-1)}||_F^2 \tag{3}$$

subject to $Z_x^{(t)}, V_x^{(t)}, Z_x^{(t-1)}, V_x^{(t-1)} \geq 0$, where $||\cdot||_F$ denotes the Frobenius norm, with $||X||_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n}|X_{ij}|^2}$. Notice that the products $Z_x^{(t)}V_x^{(t)}$ and $Z_x^{(t-1)}V_x^{(t-1)}$ are the factorized matrices $\hat{X}^{(t)}$ and $\hat{X}^{(t-1)}$, respectively. To capture the preference dynamics between the time periods t and $t-1$ based on the target interaction matrices $X^{(t)}$ and $X^{(t-1)}$, we introduce a transition matrix $T_x^{(t)} \in \mathbb{R}^{d\times d}$ between the user latent matrices V_x^t and V_x^{t-1} in the time periods t and $t-1$. The transition matrix $T_x^{(t)}$ captures how much the current/present user latent matrix V_x^t can be expressed by the past user latent matrix V_x^{t-1}. According to Vaca et al. (2014), we can reformulate the NMF of the target matrix $X^{(t)}$ in Eq. (2) as the following joint minimization problem:

$$\min_{Z_x^{(t)},V_x^{(t)},T_x^{(t)}} ||X^{(t)} - Z_x^{(t)}V_x^{(t)}||_F^2 + ||X^{(t)} - Z_x^{(t)}T_x^{(t)}V_x^{(t-1)}||_F^2 \tag{4}$$

where the goal is to minimize both the low rank d approximation errors.

Similarly, based on the auxiliary interaction matrices $A^{(t)}$ and $A^{(t-1)}$ in the periods t and $t-1$ we define a transition matrix $T_a^{(t)}$ which denotes how much the current/present user latent matrix V_a^t can be expressed by the past user latent matrix $V_a^{(t-1)}$. In our setting, we assume that $T_a^{(t)} \neq T_x^{(t)}$, as the interactions types in $X^{(t)}$ and $A^{(t)}$ do not necessarily evolve in the same way. The NMF of the auxiliary matrix $A^{(t)}$ can be rewritten as the following *joint* minimization problem:

$$\min_{Z_a^{(t)},V_a^{(t)},T_a^{(t)}} ||A^{(t)} - Z_a^{(t)}V_a^{(t)}||_F^2 + ||A^{(t)} - Z_a^{(t)}T_a^{(t)}V_a^{(t-1)}||_F^2 \tag{5}$$

To consider both the user preference dynamics and the multi-modal user-item interactions, we combine Eqs. (4) and (5), which results in the following *joint objective function* with respect to variables/matrices $Z_x^{(t)}$, $V_x^{(t)}$, $T_x^{(t)}$, $Z_a^{(t)}$, $V_a^{(t)}$ and $T_a^{(t)}$:

$$\begin{aligned}
&\min_{Z_x^{(t)},V_x^{(t)},T_x^{(t)},Z_a^{(t)},V_a^{(t)},T_a^{(t)}} \mathcal{L} \\
&= ||X^{(t)} - Z_x^{(t)}V_x^{(t)}||_F^2 + ||X^{(t)} - Z_x^{(t)}T_x^{(t)}V_x^{(t-1)}||_F^2 \\
&\quad + ||A^{(t)} - Z_a^{(t)}V_a^{(t)}||_F^2 + ||A(t) - Z_a^{(t)}T_a^{(t)}V_a^{(t-1)}||_F^2 \\
&\quad + \lambda(||T_x^{(t)} - I||_F^2 + ||T_a^{(t)} - I||_F^2) \\
&\quad + \beta(||Z_x^{(t)}||_1 + ||V_x^{(t)}||_1 + ||Z_a^{(t)}||_1 + ||V_a^{(t)}||_1) \\
&\quad + \gamma(||T_x^{(t)}||_1 + ||T_a^{(t)}||_1)
\end{aligned} \tag{6}$$

where $I \in \mathbb{R}^{d\times d}$ is the identity matrix, and $||\cdot||_1$ denotes the l_1-norm, with $||X|_1 = \max_{1\leq j\leq n}\sum_{i=1}^{m}|X_{ij}|$. In the joint objective function of Eq. (6) the first four terms denote the low rank d approximation errors based on Eqs. (4) and (5). The fifth term is a temporal regularization based on the transition matrices $T_x^{(t)}$ and $T_a^{(t)}$, where the temporal regularizer λ controls the transition between the past and user preferences, that is, higher values of λ make the model more biased to the past preferences for both the target and auxiliary interactions. The sixth term consists of the l_1-norm regularization, which forces the factor matrices $Z_x^{(t)}$, $V_x^{(t)}$, $Z_a^{(t)}$ and $V_a^{(t)}$ to be sparse (Blei, Ng, & Jordan, 2003; Mairal, Bach, Ponce, & Sapiro, 2010), with a regularizer parameter β that controls the impact on the joint objective function $\mathcal{L}$. The last term forces also the transition matrices $T_x^{(t)}$ and $T_a^{(t)}$ to be sparse, as the transition of users' multimodal preferences between the past and current periods change gradually over time, with γ being the respective regularization parameter. For simplicity, we set $\lambda = \beta = \gamma$ in our implementation.

4.2. Optimization algorithm

As the joint objective function $\mathcal{L}$ is not convex with respect to the six variables/matrices, we propose an alternating optimization algorithm based on the strategy of multiplicative update rules (Jialu et al., 2013; Lee & Seung, 1999), where we update one variable, while keeping the remaining five fixed. Using the Karush–Kuhn Tucker (KKT) conditions (Kuhn & Tucker, 1951) we have:

$$Z_x^{(t)} \geq 0, \quad V_x^{(t)} \geq 0, \quad T_x^{(t)} \geq 0, \quad Z_a^{(t)} \geq 0, \quad V_a^{(t)} \geq 0, \quad T_a^{(t)} \geq 0 \tag{7}$$

$$\begin{aligned}
&\nabla_{Zx^{(t)}}\mathcal{L} \geq 0, \quad \nabla_{Vx^{(t)}}\mathcal{L} \geq 0, \quad \nabla_{Tx^{(t)}}\mathcal{L} \geq 0, \\
&\nabla_{Za^{(t)}}\mathcal{L} \geq 0, \quad \nabla_{Va^{(t)}}\mathcal{L} \geq 0, \quad \nabla_{Ta^{(t)}}\mathcal{L} \geq 0
\end{aligned} \tag{8}$$

$$\begin{aligned}
&Z_x^{(t)} \odot \nabla_{Zx^{(t)}}\mathcal{L} = 0, \quad V_x^{(t)} \odot \nabla_{Vx^{(t)}}\mathcal{L} = 0, \quad T_x^{(t)} \odot \nabla_{Tx^{(t)}}\mathcal{L} = 0, \\
&Z_a^{(t)} \odot \nabla_{Za^{(t)}}\mathcal{L} = 0, \quad V_a^{(t)} \odot \nabla_{Va^{(t)}}\mathcal{L} = 0, \quad T_a^{(t)} \odot \nabla_{Ta^{(t)}}\mathcal{L} = 0
\end{aligned} \tag{9}$$

where $\odot$ denotes the element-wise product. The gradients of the joint objective function $\mathcal{L}$ in Eq. (6) with respect to each variable are equivalent to:

$$\begin{aligned}
\nabla_{Zx^{(t)}}\mathcal{L} =& Z_x^{(t)}(V_x^{(t)}{V_x^{(t)}}^T + T_x^{(t)}V_x^{(t-1)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T) \\
&- (X^{(t)}{V_x^{(t)}}^T + X^{(t)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T - \beta)
\end{aligned} \tag{10}$$

$$\nabla_{Vx^{(t)}}\mathcal{L} = {Z_x^{(t)}}^T Z_x^{(t)}V_x^{(t)} - ({Z_x^{(t)}}^T X^{(t)} - \beta) \tag{11}$$

$$\begin{aligned}
\nabla_{Tx^{(t)}}\mathcal{L} =& V_x^{(t-1)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T{Z_x^{(t)}}^T Z_x^{(t)} + \lambda {T_x^{(t)}}^T \\
&- (V_x^{(t-1)}{X^{(t)}}^T Z_x^{(t)} + \lambda I - \gamma)
\end{aligned} \tag{12}$$

$$\begin{aligned}
\nabla_{Za^{(t)}}\mathcal{L} =& Z_a^{(t)}(V_a^{(t)}{V_a^{(t)}}^T + T_a^{(t)}V_a^{(t-1)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T) \\
&- (A^{(t)}{V_a^{(t)}}^T + A^{(t)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T - \beta)
\end{aligned} \tag{13}$$

$$\nabla_{Va^{(t)}}\mathcal{L} = {Z_a^{(t)}}^T Z_a^{(t)}V_a^{(t)} - ({Z_a^{(t)}}^T A^{(t)} - \beta) \tag{14}$$

$$\begin{aligned}
\nabla_{Ta^{(t)}}\mathcal{L} =& V_a^{(t-1)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T{Z_a^{(t)}}^T Z_a^{(t)} + \lambda {T_a^{(t)}}^T \\
&- (V_a^{(t-1)}{A^{(t)}}^T Z_a^{(t)} + \lambda I - \gamma)
\end{aligned} \tag{15}$$

Based on the gradients in Eqs. (10)–(15), by solving with respect to each variable in the conditions of Eq. (9), the following updating rules are derived:

$$Z_x^{(t)} \leftarrow Z_x^{(t)} \odot \frac{X^{(t)}{V_x^{(t)}}^T + X^{(t)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T - \beta}{Z_x^{(t)}(V_x^{(t)}{V_x^{(t)}}^T + T_x^{(t)}V_x^{(t-1)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T)} \tag{16}$$

$$V_x^{(t)} \leftarrow V_x^{(t)} \odot \frac{{Z_x^{(t)}}^T X^{(t)} - \beta}{{Z_x^{(t)}}^T Z_x^{(t)} T_x^{(t)}} \tag{17}$$

$$T_x^{(t)} \leftarrow T_x^{(t)} \odot \frac{V_x^{(t-1)}{X^{(t)}}^T Z_x^{(t)} + \lambda I - \gamma}{V_x^{(t-1)}{V_x^{(t-1)}}^T{T_x^{(t)}}^T{Z_x^{(t)}}^T Z_x^{(t)} + \lambda {T_x^{(t)}}^T} \tag{18}$$

$$Z_a^{(t)} \leftarrow Z_a^{(t)} \odot \frac{A^{(t)}{V_a^{(t)}}^T + A^{(t)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T - \beta}{Z_a^{(t)}(V_a^{(t)}{V_a^{(t)}}^T + T_a^{(t)}V_a^{(t-1)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T)} \tag{19}$$

$$V_a^{(t)} \leftarrow V_a^{(t)} \odot \frac{{Z_a^{(t)}}^T A^{(t)} - \beta}{{Z_a^{(t)}}^T Z_a^{(t)} T_a^{(t)}} \tag{20}$$

$$T_a^{(t)} \leftarrow T_a^{(t)} \odot \frac{V_a^{(t-1)}{A^{(t)}}^T Z_a^{(t)} + \lambda I - \gamma}{V_a^{(t-1)}{V_a^{(t-1)}}^T{T_a^{(t)}}^T{Z_a^{(t)}}^T Z_a^{(t)} + \lambda {T_a^{(t)}}^T} \tag{21}$$

Algorithm 1 presents the proposed Temporal Collective Matrix Factorization (TCMF) method. In line 2, we initialize the six

Algorithm 1: The TCMF algorithm.

Input: $X^{(t)}$, $X^{(t-1)}$, $A^{(t)}$, $A^{(t-1)}$, $\{\lambda, \beta, \gamma, \epsilon, d\}$, $maxIter$
Output: $\hat{X}^{(t)}$

1 $\theta' \leftarrow$ maxInit, $\theta \leftarrow \frac{\theta'}{2}$
2 Initialize $Z_x^{(t)}$, $V_x^{(t)}$, $T_x^{(t)}$, $Z_a^{(t)}$, $V_a^{(t)}$, $T_a^{(t)}$
3 Compute $V_x^{(t-1)}$ and $V_a^{(t-1)}$
4 **while** $(abs(\theta' - \theta) > \epsilon) \vee (iter < maxIter)$ **do**
5   Update $Z_x^{(t)}$ (Equation (16))
6   Update $V_x^{(t)}$ (Equation (17))
7   Update $T_x^{(t)}$ (Equation (18))
8   Update $Z_a^{(t)}$ (Equation (19))
9   Update $V_a^{(t)}$ (Equation (20))
10   Update $T_a^{(t)}$ (Equation (21))
11   Compute $\mathcal{L}$ based on the updated $Z_x^{(t)}$, $V_x^{(t)}$, $T_x^{(t)}$, $Z_a^{(t)}$, $V_a^{(t)}$, $T_a^{(t)}$
12   (Equation (6))
13     $\theta' \leftarrow \theta$
14     $\theta \leftarrow \mathcal{L}$
15     $iter \leftarrow iter + 1$
    end
16 $\hat{X}^{(t)} = Z_x^{(t)} V_x^{(t)}$

variables as follows: the factor matrices $Z_x^{(t)}$, $V_x^{(t)}$, $Z_a^{(t)}$ and $V_a^{(t)}$ are initialized by random (sparse) matrices, with $Z_x^{(t)}$, $V_x^{(t)}$, $Z_a^{(t)}$ and $V_a^{(t)} \geq 0$. The transition matrices are initialized by setting $T_x^{(t)}, T_a^{(t)} \leftarrow \lambda I$. In line 3, we compute the user factor matrices $V_x^{(t-1)}$ and $V_a^{(t-1)}$ by applying NMF to $X^{(t-1)}$ and $A^{(t-1)}$, that is, the target and the auxiliary matrices for the past time period $t-1$. Our iterative optimization algorithm starts in lines 4–14, where we update the six variables/matrices based on Eqs. (16)–(21), respectively. The order of the matrices' updating rules does not influence the solutions, as the matrices take similar values after a few iterations (Jialu et al., 2013; Lee & Seung, 1999). Then, in line 11 we calculate the joint objective function $\mathcal{L}$ in Eq. (6) based on the updated matrices. Notice that in each iteration, the optimization algorithm minimizes the joint objective function.[2] After computing the updated $\mathcal{L}$, in line 4 we examine if the algorithm has converged, that is, (i) we calculate the absolute difference between the current and the previous iteration and if the difference is larger than a predefined convergence threshold ϵ we continue the optimization algorithm or (ii) if we have reached the maximum number of iterations *maxIter*. In our implementation, we fix $\epsilon = 10^{-6}$ and $maxIter = 10^5$. After the optimization algorithm finishes, in line 15 we compute the factorized matrix $\hat{X}^{(t)}$, which contains the final recommendations.

5. Experimental evaluation

5.1. Data sets

In our experiments we use two real-world data sets from the Stanford Network Analysis Project (SNAP), *Fine Food*[3] and *Movies*[4] (McAuley & Leskovec, 2013). Their main characteristics are presented in Table 2. Both data sets contain bimodal user-item interactions, that is, users' comments and ratings on Amazon's products. The data sets *Fine Food* and *Movies* span at least 13 years, having users' bimodal interactions during 156 and 180 months, respectively. Fig. 1 presents the number of ratings and comments per month in the evaluation data sets. We observe that users choose to rate products more often than comment. Hence, in our experiments users' ratings are considered as the target interaction type, while the interaction type of comments is considered as the auxiliary one. In addition, we observe that the monthly numbers of both interactions types are significantly increased after a few months.

Table 2
Evaluation data sets.

Data sets	Fine Food	Movies
Users	256,059	889,176
Items	74,258	253,059
Ratings	568,454	7,911,684
Comments	264,628	5,059,132
Time span	Oct.1999–Oct.2012	Aug.1997–Oct.2012

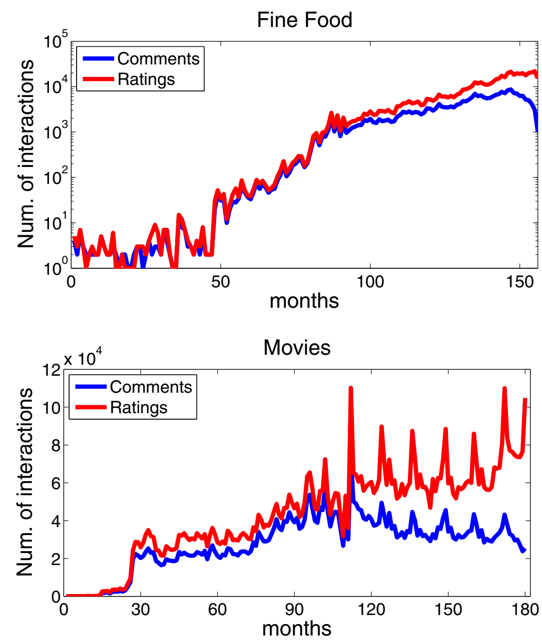

Fig. 1. Number of user-item interactions for the 156 and 180 months in *Fine Food* and *Movies*, respectively.

5.2. Evaluation protocol

The evaluation data sets are split into training and test months, that is, given a test month over the time span, we train the examined methods on all the past months. For example, if the 10th month is the test set, we train the examined methods on the past 9 months. For each test month, we average the results over all the users.

We evaluate the performance in terms of the AUC metric, which measures the area under the Receiver Operating Characteristic (ROC) curve - True Positives versus False Positives. AUC is the proportion of correctly ordered items in the ranked list (Liu et al., 2013). A pair of items is *discordant*, if the item at the higher rank position is irrelevant while the result at the lower rank position is relevant. AUC is defined as follows:

$$AUC = 1 - \frac{N^{\mp}}{N^{+} \times N^{-}} \quad (22)$$

[2] As we considered the optimality condition in Eq. (9) to generate the updating rules, the proof that the optimization algorithm minimizes the joint objective function in each iteration, and thus converges, is similar to the study at Lee and Seung (1999).
[3] http://snap.stanford.edu/data/web-FineFoods.html.
[4] http://snap.stanford.edu/data/web-Movies.html.

where N^+ and N^- denote the total number of *relevant* and *irrelevant* items of a user, respectively, and $N^{\mp}$ denotes the number of *discordant* item pairs. Higher values of AUC express a high recommendation accuracy.

5.3. Compared methods

In our experiments, we evaluate the performance of the following methods:

- *timeSVD++* (Koren, 2009): is a baseline method which factorizes a single user-item interaction matrix, by taking into account the users' preference dynamics. timeSVD++ extends matrix factorization by introducing time dependent user/item biases and user latent factors to cope with the dynamics of user preferences. However, timeSVD++ ignores the auxiliary interaction types, when generating the recommendations in a test month.
- *Temporal Matrix Factorization (TMF)*[5] (Zhang et al., 2014a): models users' preference dynamics by learning a transition matrix for each user latent vectors between two time periods. In TMF, the transition matrix considers the time-invariant pattern of the preference dynamics. TMF exploits only a single type of user-item interactions.
- *Bayesian Temporal Matrix Factorization* (BTMF)[5] (Zhang et al., 2014a): extends the TMF method by introducing priors for the hyperparameters which capture the conditional distributions of users, items and a single type of interactions. As TMF, different types of user-item interactions are not captured in the BTMF method.
- *Local Collective Embeddings* (LCE)[6] (Saveski & Mantrach, 2014): is a collective matrix factorization strategy that co-factorizes multimodal user-item interactions. In relation to our approach, LCE assumes that the multimodal user-interaction matrices X and A share a common item factor matrix Z, and then LCE tries to minimize the following objective function:

$$\min_{Z,V_x,V_a} ||X - ZV_x|_F^2 + ||A - ZV_a||_F^2 + \beta(||Z||_F^2 + ||V_x||_F^2 + ||V_a||_F^2) \tag{23}$$

  where V_x and V_z are the respective user factor matrices based on X and A, and the third term is a L_2-norm regularizer on Z, V_x and V_a. LCE does not capture the transition of users' preference over time, when minimizing the objective function in Eq. (23). Hence, LCE is trained on the past months as the temporal models timeSVD++, TMF and BTMF, but does not consider the time dimension, which means that LCE ignores users' preference dynamics between the months.
- *TCMF*: is the proposed Temporal Collective Matrix Factorization method, which considers both the users' preference dynamics and multimodal user-item interactions.

5.4. Impact of the temporal regularizer λ

In the proposed TCMF method, we balance the influence of the past multimodal interactions based on the temporal regularizer λ, where higher values of λ bias more the model to the past preferences. In Eq. (2) we evaluate the performance of TCMF in terms of AUC, by averaging the results over all the test months. As aforementioned in Section 5.2, high values of AUC are translated into higher recommendation accuracy. We observe that a conservative selection of $\lambda = 1$ is required for both data sets, as higher values of λ make TCMF more biased to the past preferences and as a consequence TCMF ignores users' more recent selections, while lower values of λ make TCMF forget users' past preferences.

[5] http://www.sfu.ca/~chenyiz/content/TMF/code.rar.

[6] http://www.time.mk/msaveski/code.htm.

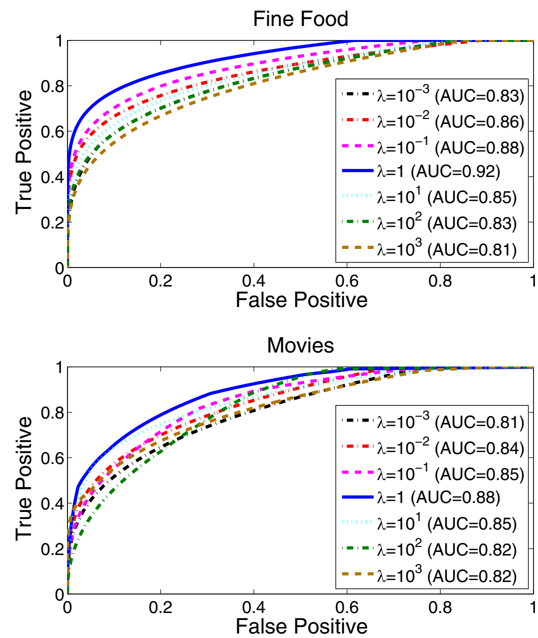

Fig. 2. Effect of temporal regularizer λ on AUC. The results are averaged over the test months.

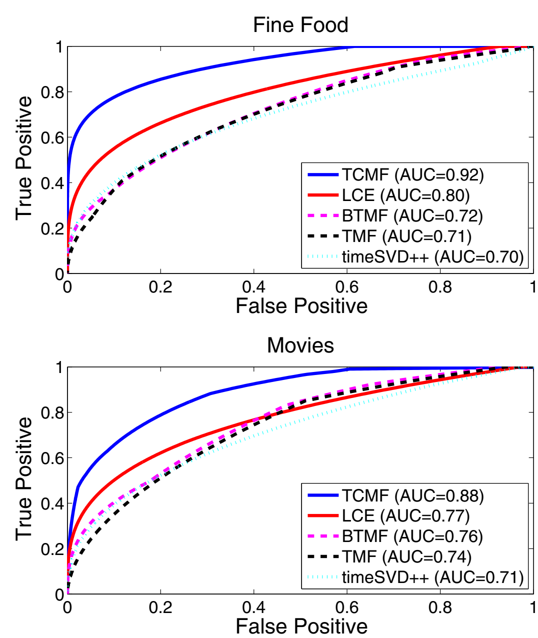

Fig. 3. Comparison of TCMF with baselines in terms of AUC.

5.5. Comparison with state-of-the-art

In the next set of experiments, we compare the proposed TCMF method with the temporal methods timeSVD++, TMF and BTMF. We also compare with the LCE method which captures the bimodal interactions, but ignores the preference dynamics. As we can observe from Fig. 3 timeSVD++, TMF and BTMF capture the preference dynamics over the time span, but do not exploit the bimodal user-item interactions, thus having a glass ceiling on the recommendation accuracy due to the data sparsity problem. Meanwhile, LCE achieves slightly higher recommendation accuracy than the temporal methods, by exploiting the bimodal user-item interaction in the evaluation data sets. However, LCE has also a glass ceiling on the recommendation accuracy, as it does not consider users' pref-

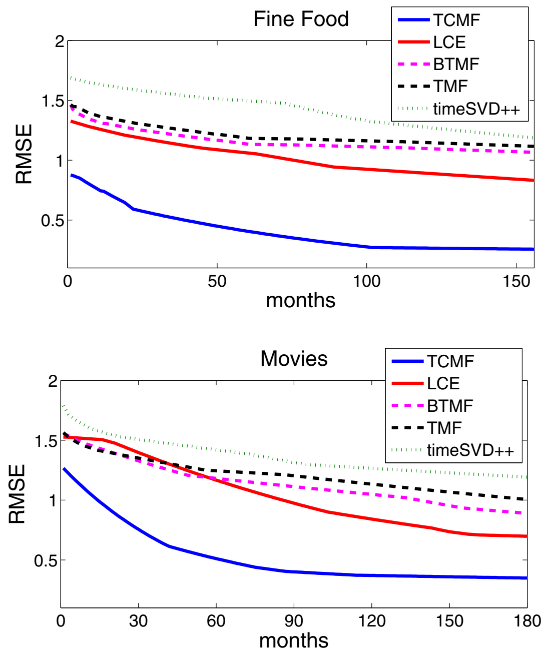

Fig. 4. Comparison of TCMF with baselines in terms of RMSE.

erence dynamics over the time span. The proposed TCMF method constantly outperforms all the competitive methods. This occurs because TCMF exploits the bimodal user-item interactions and captures the users' preference dynamics over the data sets' evolution, both important factors when generating the recommendations.

5.6. Performance on rating prediction

To demonstrate the superiority of the proposed TCMF method over the competitors in different recommendation tasks, in this set of experiments we focus on the rating prediction problem. Notice that main difference between the item recommendation and rating prediction problems is that methods that try to predict the values of the ratings may not provide the best recommendation lists to the end-user in many cases. For example, if all the low-ranked ratings are predicted very accurately, but significant errors are made on the higher-ranked ratings, the resulting solution will not provide a high-quality recommendation list to the end-user and vice versa (Aggarwal, 2016). A common evaluation metric in rating prediction is the *Root Mean Square Error* (RMSE) (Levinson, 1947). Given r ratings of a test month, the ground truth ratings in X and the predicted ratings in $\hat{X}$, RMSE is defined as follows:

$$\text{RMSE} = \sqrt{\frac{\sum_{i,j}(X_{i,j} - \hat{X}_{i,j})^2}{r}} \tag{24}$$

where lower values of RMSE indicate higher performance in the rating prediction problem. Figs. 4 shows the effect on RMSE for the examined methods in Fine Food and Movies over the datasets' time span. We can observe that all methods improve their learning ability because in each time period the training set is augmented. For example, the temporal methods timeSVD++, TMF and BTMF decrease the RMSE metric until a month, close to the 100th month and the 90th month in Fine Food and Movies, respectively, and then there is a slight improvement over the rest time span. For all test months over the datasets' time span the proposed TCMF method keeps the rating prediction relatively high, by beating the competitors in the rating prediction problem, as well.

5.7. Discussion

Summarizing our results, we observe that capturing users' preference dynamics and exploiting multimodal user-item interactions, are both crucial factors in recommender systems. Methods that consider users' preference dynamics can capture how users shift their preferences over time. However, such temporal methods e.g., timeSVD++, TMF and BTMF, consider only one type of user-item interactions, thus not handling the data sparsity problem. This explains their limited accuracy over the datasets' time span in the experimental evaluation. On the other hand, LCE is a recommendation algorithm that can exploit multimodal user-item interactions, but does not capture preference dynamics. Our results clearly show that the proposed TCMF method beats all its competitors, in the top-N recommendation and rating prediction problems. This occurs because TCMF is the only method that considers both users' preference dynamics of multimodal user-item interactions.

6. Conclusion

In this article we presented TCMF, an efficient social media recommendation algorithm to capture the preference dynamics in the multimodal interactions that users perform over the time evolution. The proposed TCMF method achieves high recommendation accuracy over the datasets' time span in different recommendation tasks, as competitive approaches either consider users' preference dynamics or their evolution over time, but not both of them. Interesting future directions are:

- Consider the social relationships when generating the recommendations, assuming that users tend to trust the selections of their social friends (Guo, Zhang, & Yorke-Smith, 2015; Jamali & Ester, 2010).
- Account for the fact that user preferences may be updated based on the preferences of their friends over time, while their social relationships may change over time as well (Liu et al., 2013). As future work, we plan to extend our model to capture the time dimension of the social relationships.
- Develop a cross-domain recommendation algorithm that will exploit user preferences based on multimodal user-item interactions in different domains (Gao et al., 2013; Mirbakhsh & Ling, 2015; Rafailidis & Crestani, 2016c).
- Examine different collaborative ranking models that perform push at the top of the list, considering that what matters is the ranking performance at the top of the list, the recommendations that the user will actually see (Christakopoulou & Banerjee, 2015; Rafailidis & Crestani, 2016a, b).

References

Aggarwal, C. C. (2016). *Recommender systems - The textbook*. Springer.
Ahn, H. J. (2008). A new similarity measure for collaborative filtering to alleviate the new user cold-starting problem. *Information Sciences, 178*(1), 37–51.
Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research, 3*, 993–1022.
Bobadilla, J., Ortega, F., Hernando, A., & Bernal, J. (2012). A collaborative filtering approach to mitigate the new user cold start problem. *Knowledge-Based Systems, 26*, 225–238.
Bouchard, G., Yin, D., & Guo, S. (2013). Convex collective matrix factorization. In *Proceedings of the 16th international conference on artificial intelligence & statistics (aistats)* (pp. 144–152).
Caicedo, J. C., BenAbdallah, J., González, F. A., & Nasraoui, O. (2012). Multimodal representation, indexing, automated annotation and retrieval of image collections via non-negative matrix factorization. *Neurocomputing, 76*(1), 50–60.
Cho, E., Myers, S. A., & Leskovec, J. (2011). Friendship and mobility: User movement in location-based social networks. In *Proceedings of the 17th ACM international conference on knowledge discovery & data mining (kdd)* (pp. 1082–1090).
Christakopoulou, K., & Banerjee, A. (2015). Collaborative ranking with a push at the top. In *Proceedings of the 24th international conference on world wide web, WWW 2015, florence, italy, may 18–22, 2015* (pp. 205–215).

Da Costa, A., Domingues, M., Rezende, S., & Manzato, M. (2014). Improving personalized ranking in recommender systems with multimodal interactions. In *Proceedings IEEE/WIC/ACM international joint conferences on web intelligence (WI) & intelligent agent technologies (IAT): 1* (pp. 198–204).

Da Costa, A., & Manzato, M. (2014). Multimodal interactions in recommender systems: An ensembling approach. In *Proceedings of the brazilian conference on intelligent systems (BRACIS)* (pp. 67–72).

Fang, Y., & Si, L. (2011). Matrix co-factorization for recommendation with rich side information and implicit feedback. In *Proceedings of the 2nd international workshop on information heterogeneity & fusion in recommender systems (HetRec)* (pp. 65–69).

Gao, S., Luo, H., Chen, D., Li, S., Gallinari, P., & Guo, J. (2013). Cross-domain recommendation via cluster-level latent factor model. In *Machine learning and knowledge discovery in databases - European conference, ECML PKDD 2013, prague, czech republic, september 23–27, 2013, proceedings, part II* (pp. 161–176).

Guo, G., Zhang, J., & Yorke-Smith, N. (2015). Trustsvd: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings. *AAAI*.

Herlocker, J. L., Konstan, J. A., Terveen, L. G., & Riedl, J. T. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems, 22*(1), 5–53.

Huang, Y.-J., Xiang, E. W., & Pan, R. (2012). Constrained collective matrix factorization. In *Proceedings of the 6th ACM conference on recommender systems (RecSys)* (pp. 237–240).

Jamali, M., & Ester, M. (2010). A matrix factorization technique with trust propagation for recommendation in social networks. In *RecSys* (pp. 135–142).

Jialu, L., Chi, W., Gao, J., & Jiawei, H. (2013). Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the 13th SIAM international conference on data mining (SDM)* (pp. 252–260).

Kim, B. S., Kim, H., Lee, J., & Lee, J.-H. (2014). Improving a recommender system by collective matrix factorization with tag information. In *Proceedings 7th international conference on soft computing & intelligent systems (SCIS)* (pp. 980–984).

Koren, Y. (2009). Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM international conference on knowledge discovery & data mining (KDD)* (pp. 447–456). ACM.

Kuhn, H. W., & Tucker, A. W. (1951). Nonlinear programming. In *Proceedings of the 2nd Berkeley symposium on mathematical statistics & probability* (pp. 481–492). Berkeley, California.

Lathia, N., Hailes, S., Capra, L., & Amatriain, X. (2010). Temporal diversity in recommender systems. In *Proceedings of the 33rd ACM international conference on research & development in information retrieval (SIGIR)* (pp. 210–217).

Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature, 401*, 788–791.

Lee, D. D., & Seung, H. S. (2000). Algorithms for non-negative matrix factorization. In *Proceedings of the annual conference on neural information processing systems (NIPS), denver, co, usa* (pp. 556–562).

Levinson, N. (1947). The wiener root mean square error criterion in filter design and prediction. *Journal of Mathematics & Physics, 25*, 261–278.

Li, Z., Liu, J., Zhu, X., & Lu, H. (2010). Multi-modal multi-correlation person-centric news retrieval. In *Proceedings of the 19th ACM international conference on information & knowledge management (CIKM)* (pp. 179–188).

Liu, N. N., He, L., & Zhao, M. (2013). Social temporal collaborative ranking for context aware movie recommendation. *ACM Transactions on Intelligent Systems & Technology, 4*(1), 15:1–15:26.

Mairal, J., Bach, F., Ponce, J., & Sapiro, G. (2010). Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research, 11*, 19–60.

McAuley, J. J., & Leskovec, J. (2013). From amateurs to connoisseurs: Modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on world wide web (WWW)* (pp. 897–908).

Mirbakhsh, N., & Ling, C. X. (2015). Improving top-n recommendation for cold-start users via cross-domain information. *TKDD, 9*(4), 33.

Narita, A., Hayashi, K., Tomioka, R., & Kashima, H. (2011). Tensor factorization using auxiliary information. In *Proceedings of the European conference on machine learning & knowledge discovery in databases (ECML/PKDD), part ii* (pp. 501–516).

Porteous, I., Asuncion, A., & Welling, M. (2010). Bayesian matrix factorization with side information and dirichlet process mixtures. *Aaai conference on artificial intelligence*.

Rafailidis, D., & Crestani, F. (2016). Collaborative ranking with social relationships for top-n recommendations. In *Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval, SIGIR 2016, Pisa, Italy, july 17–21, 2016* (pp. 785–788).

Rafailidis, D., & Crestani, F. (2016). Joint collaborative ranking with social relationships in top-n recommendation. In *Proceedings of the 25th ACM international on conference on information and knowledge management, CIKM 2016, Indianapolis, IN, USA, october 24–28, 2016* (pp. 1393–1402).

Rafailidis, D., & Crestani, F. (2016). Top-n recommendation via joint cross-domain user clustering and similarity learning. In *Machine learning and knowledge discovery in databases - European conference, ECML PKDD 2016, Riva del Garda, Italy, september 19–23, 2016, proceedings, part II* (pp. 426–441).

Rafailidis, D., & Nanopoulos, A. (2014). Modeling the dynamics of user preferences in coupled tensor factorization. In *Proceedings of the 8th acm conference on recommender systems (RecSys)* (pp. 321–324).

Rafailidis, D., & Nanopoulos, A. (2015). Repeat consumption recommendation based on users preference dynamics and side information. In *Proceedings of the 24th international conference on world wide web companion (WWWCompanion)* (pp. 99–100).

Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on world wide web (WWW)* (pp. 285–295).

Saveski, M., & Mantrach, A. (2014). Item cold-start recommendations: Learning local collective embeddings. In *Proceedings of the 8th acm conference on recommender systems (RecSys)* (pp. 89–96).

Singh, A. P., & Gordon, G. J. (2008). Relational learning via collective matrix factorization. In *Proceedings of the 14th acm international conference on knowledge discovery & data mining (KDD)* (pp. 650–658).

Tso, K. H. L., Marinho, L. B., & Schmidt-Thieme, L. (2008). Tag-aware recommender systems by fusion of collaborative filtering algorithms. In *Proceedings of the acm symposium on applied computing (SAC)* (pp. 1995–1999). New York, NY

Vaca, C. K., Mantrach, A., Jaimes, A., & Saerens, M. (2014). A time-based collective factorization for topic discovery and monitoring in news. In *Proceedings of the 23rd international conference on world wide web (WWW)* (pp. 527–538).

Xiong, L., Chen, X., Huang, T.-K., Schneider, J., & Carbonell, J. G. (2010). Temporal collaborative filtering with bayesian probabilistic tensor factorization. In *Proceedings of the 10th SIAM international conference on data mining (SDM)* (pp. 211–222).

Yin, H., Cui, B., Chen, L., Hu, Z., & Zhou, X. (2015). Dynamic user modeling in social media systems. *ACM Transactions on Information Systems, 33*(3), 10:1–10:44.

Yin, Z., Gupta, M., Weninger, T., & Han, J. (2010). A unified framework for link recommendation using random walks. In *Proceedings of the 2nd IEEE international conference on advances in social networks analysis & mining (ASONAM)* (pp. 152–159).

Zhang, C., Wang, K., Yu, H., Sun, J., & Lim, E. (2014). Latent factor transition for dynamic collaborative filtering. In *Proceedings of the 14th SIAM international conference on data mining (SDM)* (pp. 452–460).

Zhang, H., Li, Z., Chen, Y., Zhang, X., & Wang, S. (2014). Exploit latent dirichlet allocation for one-class collaborative filtering. In *Proceedings of the 23rd ACM international conference on information & knowledge management(CIKM)* (pp. 1991–1994).

Zhou, T., Shan, H., Banerjee, A., & Sapiro, G. (2012). Kernelized probabilistic matrix factorization: Exploiting graphs and side information. In *Proceedings of the 12th SIAM international conference on data mining (SDM)* (pp. 403–414).