


Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Information Processing and Management

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

A dual-perspective latent factor model for group-aware social event recommendation



Yogesh Jhamb, Yi Fang*

Department of Computer Engineering, Santa Clara University, 500 El Camino Real, Santa Clara, California 95053, USA

ARTICLE INFO

Article history:

Received 1 July 2016

Revised 24 November 2016

Accepted 3 January 2017

Keywords:

Event-based social networks

Social event recommendation

Latent factor models

ABSTRACT

Event-based social networks (EBSNs) have experienced increased popularity and rapid growth. Due to the huge volume of events available in EBSNs, event recommendation becomes essential for users to find suitable events to attend. Different from classic recommendation scenarios (e.g., movies and books), a large majority of EBSN users join groups unified by a common interest, and events are organized by groups. In this paper, we propose a dual-perspective latent factor model for group-aware event recommendation by using two kinds of latent factors to model the dual effect of groups: one from the user-oriented perspective (e.g., topics of interest) and another from the event-oriented perspective (e.g., event planning and organization). Pairwise learning is used to exploit unobserved RSVPs by modeling the individual probability of preference via Logistic and Probit sigmoid functions. These latent group factors alleviate the cold-start problems, which are pervasive in event recommendation because events published in EBSNs are always in the future and many of them have little or no trace of historical attendance. The proposed model is flexible and we further incorporate additional contextual information such as event venue, event popularity, temporal influence and geographical distance. We conduct a comprehensive set of experiments on four datasets from *Meetup* in both regular and cold-start settings. The results demonstrate that the proposed approach yields substantial improvement over the state-of-the-art baselines by utilizing the dual latent factors of groups.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

Event recommender systems have recently gained prevalence with the advent of *Event-Based Social Networks (EBSNs)*. EBSNs allow like-minded people to gather together and socialize on a wide range of topics. Among all the elements in EBSNs, events are the most significant one, which bridges the gap of online and offline interactions. As of December 2015, *Meetup*,¹ one of the largest EBSNs today, has over 24 million members, with approximately 200,000 groups in 181 countries. There are approximately 500,000 events organized every month on *Meetup*. The sheer volume of available events, especially in large cities, often undermines the users' ability to find the ones that best match their interests. Consequently, personalized event recommendation is essential for overcoming such an information overload.

* Corresponding author.

E-mail address: yfang@scu.edu (Y. Fang).¹ <http://www.meetup.com>.

Users indicate their interest to attend an event by responding to a RSVP² for the event. *Meetup* generates over 3 million RSVPs every month. The RSVP indicates a user's preference on an event, and it allows future events to be recommended to the user. At first glance, event recommendation is the same as recommending any other kind of items (e.g., movies and books), with the only difference that the item here refers to an event. However, the key distinction is that EBSNs allow users to organize themselves into groups that are created based on a topic of interest. The events are hosted by groups at venues that are often in the vicinity of the local community. Such group structures are generally not available for other recommendation problems. We can view the group information from a dual perspective: user-oriented and event-oriented. The user-oriented perspective regards a group as a topic of interest so that users associated with a group are interested in the same topic with the group. On the other hand, the event-oriented perspective views a group as an organizing entity. The events organized by the same group have the same organizing style such as logistics, event planning, structure, quality of talks, etc. These two perspectives complement each other and together they form a complete view of a group. For event recommendation, an interesting question here is: how can we leverage this dual perspective of group information to provide effective event recommendation?

Moreover, optimal use of group information can largely alleviate the cold-start problems, which are pervasive in the setting of event recommendation. New events and new users are constantly emerging in EBSNs. Many events published in EBSNs have little or no trace of prior attendance because the events are always in the future and they are often short-lived. Also, as EBSNs grow rapidly, there are many new users joining without record of historical attendance. By knowing the group that organizes the new event, we can expect the organizing style of the event based on the event-oriented perspective of groups. Similarly, by looking at the groups that the new user is associated with, we can infer the interests of the user based on the user-oriented perspective of groups. Therefore, this dual perspective of groups can help address both new item and new user cold-start problems.

In EBSNs, a user may RSVP for an event in the affirmative by a positive response (“yes”), or the user may provide a negative response to an event with a RSVP as (“no”). The numbers of positive responses and negative responses are largely disproportional. Many users just ignore RSVPs if they are not interested in attending the events. Therefore, it is more desirable to treat event recommendation as the top-*N* ranking task (Kassak, Kompan, & Bielikova, 2016) than a binary rating prediction problem. On the other hand, the absence of a response does not necessarily mean that the user is not interested in the event. It may be that the user is not aware of the event, or that the user is unable to attend this event due to other conflicts. Thus, the event recommendation model needs to take into account not only the positive and negative RSVPs, but also the missing/unobserved RSVPs.

Event recommendation is much less studied in the literature than traditional recommendation tasks such as movie and book recommendations. To address the unique characteristics of event recommendation, we propose a dual-perspective group-aware latent factor model. The proposed model utilizes pairwise ranking by taking unobserved RSVPs into account. In addition to the typical user and item latent factors, two novel latent factors are used to model a group: one for its user-oriented characteristics and another for its event-oriented characteristics. The influences of the groups on the user is then modeled as the linear combination of the latent factors for the user-oriented characteristics of its groups. The experimental results show that the proposed model outperforms the state-of-the-art baselines. The results also indicate that the performance can be further improved when incorporating factors associated with event venue, event popularity, temporal influence and geographical distance. It is worth noting that while adding more features helps, the group influence drives the most performance gain and it is the focus of this work. The main contributions of this paper can be summarized as follows.

- We characterize two different perspectives of groups (user-oriented vs. event-oriented). We make use of these two complementary perspectives in event recommendation, especially for addressing the cold-start problems. To the best of our knowledge, no prior work has studied the dual-perspective of group influence on event recommendation.
- We propose a probabilistic latent factor model by incorporating two different types of latent factors to represent the user-oriented and event-oriented characteristics of groups. Pairwise learning is used to exploit unobserved RSVPs by modeling the individual probability of preference via Logistic and Probit sigmoid functions.
- The proposed model is flexible to further incorporate additional contextual information including event venue, event popularity, temporal influence and geographical distance. We incorporate these additional parameters in our model and study their impact on recommendations.
- We thoroughly evaluate our proposed approach on four datasets from *Meetup*. The results demonstrate its effectiveness compared to the state-of-the-art baselines. The source code and data will be made publicly available once the work is published.

The remainder of the paper is organized as follows. [Section 2](#) covers the related work that is relevant to our study. [Section 3](#) provides data analysis that motivates our models. [Section 4](#) introduces our event recommender models in detail, and [Section 5](#) discusses the experimental setup and the results of this study. [Section 6](#) concludes with a summary and an outline of the future work that will follow this study.

² RSVP is a French expression, which means “please respond”.

2. Related work

2.1. Event recommendation in EBSNs

Event recommendation has recently garnered increased attention with the advent of event-based social networks. A study (Liu et al., 2012) on EBSNs investigates unique and interesting characteristics of such networks by highlighting the association between the online and offline social worlds for recommending new events. Qiao et al. (2014b) later extended the work by proposing a Bayesian matrix factorization approach and employing social regularization factors inspired by user interactions in an EBSN. In Qiao et al. (2014a), a standard matrix factorization approach, which jointly models event, location, and social relation is proposed; however, they ignore content and organizer information of events. Several studies have utilized content information for event recommendation. Minkov, Charrow, Ledlie, Teller, and Jaakkola (2010) combine content-based filtering and collaborative filtering to recommend scientific seminar events. User preferences for an event were inferred based upon their preference for past events with similar content. Khrouf and Troncy (2013) leveraged an enriched content-based representation of music related events by exploiting category information from DBpedia³ about the artists associated with each event. Zhang and Wang (2015) propose a collective Bayesian Poisson factorization to jointly model user response to events, social relation, and content text.

In addition to content modeling, some recent work found that contextual information is very useful in event recommendation. Du et al. (2014) considered spatial and temporal context to predict event attendance. We have also included the geographical distance and temporal influence in our proposed model, however, the primary emphasis of our work is the dual-perspective of groups. de Macedo and Marinho (2014) conducted a large-scale analysis of several factors that impact user preferences on events in an EBSN. They observed that users tend to provide RSVPs close to the occurrence of the events. Macedo, Marinho, and Santos (2015) further proposed a context-aware approach by exploiting various contextual information including social signals based on group memberships, location signals based on the users' geographical preferences, and temporal signals derived from the users' time preferences. Chen and Sun (2016) proposed a social event recommendation method that exploits a user's social interaction relations and collaborative friendships.

An existing study tackled a related problem in EBSNs: event-based group recommendation. Zhang, Wang, and Feng (2013) exploited matrix factorization to model interactions between users and groups. By considering both explicit features (e.g., location and social features) and implicit patterns, the proposed approach demonstrated improved performance for group recommendations. This is different from our work as the objective of this study is to recommend groups that a user can join in an EBSN, whereas our work utilizes the dual-perspective of groups to recommend events to the user. A group recommender for movies is proposed based on content similarity and popularity (Pera & Ng, 2013). A recent study (Pham, Li, Cong, & Zhang, 2015) proposed a general graph-based model to solve three recommendation tasks on EBSNs in one framework, namely recommending groups to users, recommending tags to groups, and recommending events to users. The work models the rich information with a heterogeneous graph and considers the recommendation problem as a query-dependent node proximity problem. Another study (Jiang & Li, 2016) on Meetup investigated how social network, user profiles and geo-locations affect user participation when the social event is held by a single organizer. Lu et al. (2016) presented a system that extracts events from multiple data modalities and recommends events related to the user's ongoing search based on previously selected attribute values and dimensions of events being viewed.

While some recent work makes use of group memberships for event recommendation such as Macedo et al. (2015); Zhang and Wang (2015), the key distinction is that we systematically study the group influence from the dual perspective of the user and event and propose a latent factor model to formally encode this dual view.

2.2. Latent factor modeling

Latent factor models, such as matrix factorization (Koren, Bell, & Volinsky, 2009), probabilistic matrix factorization (Mnih & Salakhutdinov, 2007), and other variants (Agarwal & Chen, 2009; Bell, Koren, & Volinsky, 2007; Cao, 2015; Koren, 2010; Zhao, McAuley, & King, 2015) have demonstrated effectiveness in various recommendation tasks (Hu, Sun, & Liu, 2014; Shi, Zhao, Wang, Larson, & Hanjalic, 2012; Yao et al., 2015). Among various MF models proposed, SVD++ (Koren, 2008) is one of the most widely used models. SVD++ integrates the implicit feedback information from a user to items and the user latent factors are complemented by the latent factors of the items to which the user has provided implicit feedback.

Matrix factorization has been adapted to learn from relative pairwise preferences rather than absolute ones. One of the most effective techniques is based on Bayesian Personalized Ranking (BPR) (Rendle, Freudenthaler, Gantner, & Schmidt-Thieme, 2009), which has been shown to provide strong results in many recommendation tasks. Several extensions of BPR include pairwise interaction tensor factorization (Rendle & Schmidt-Thieme, 2010), multi-relational matrix factorization (Krohn-Grimberghe, Drumond, Freudenthaler, & Schmidt-Thieme, 2012), richer interactions among users (Pan & Chen, 2013), and non-uniformly sampled items (Gantner, Drumond, Freudenthaler, & Schmidt-Thieme, 2012). Other pairwise learning based collaborative filtering models include EigenRank (Liu & Yang, 2008) and probabilistic latent preference analysis

³ <http://wiki.dbpedia.org>.

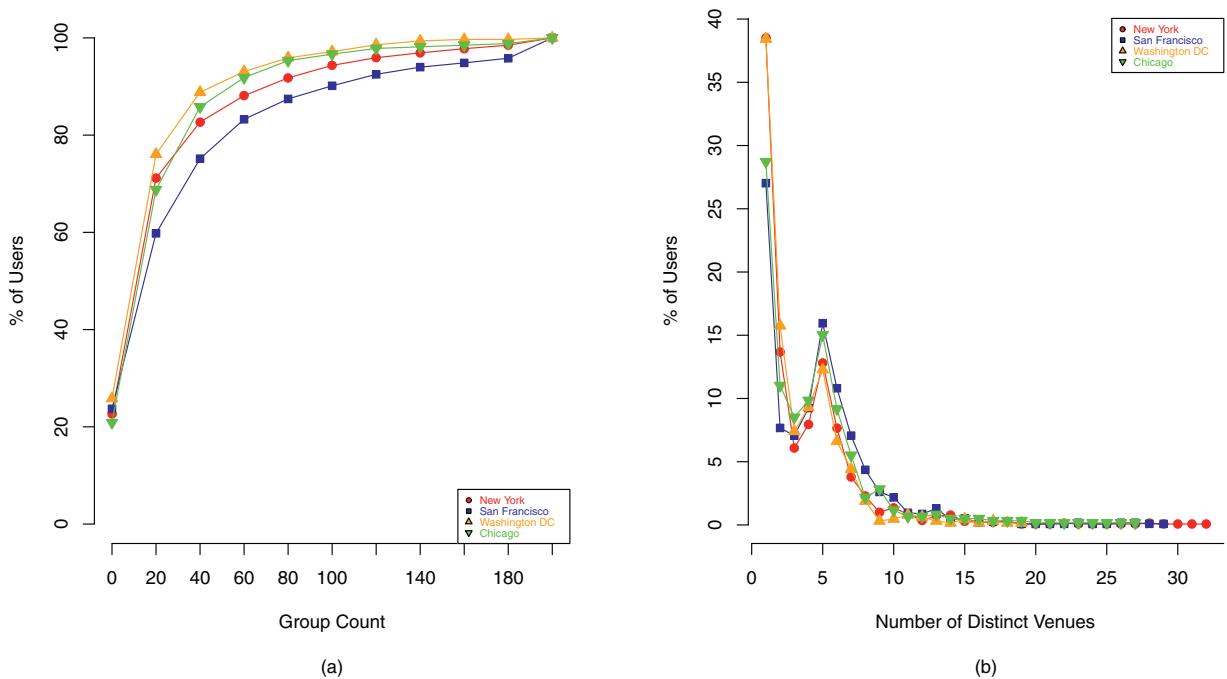


Fig. 1. User-group and user-venue relation.

(Liu, Zhao, & Yang, 2009). A pairwise ranking based geographical factorization method was recently proposed (Li, Cong, Li, Pham, & Krishnaswamy, 2015) for point-of-interest recommendation. We also utilize the pairwise ranking approach in our work, but we address the task of social event recommendation with the focus on utilizing the dual perspective of groups. Moreover, we use both logistic and Probit functions to model the pairwise preferences.

2.3. Cold-start problem

Cold-start is a prevalent problem in recommender systems as it is generally difficult for a model to handle new users and items. The cold-start problem is often alleviated by utilizing content information (Formoso, FernáNdez, Cacheda, & Carneiro, 2013). Word-based similarity methods (Pazzani & Billsus, 2007) recommend items based on textual content similarity in word vector space. Collaborative Topic Regression (CTR) couples a matrix factorization model with probabilistic topic modeling to generalize to unseen items (Wang & Blei, 2011). Collective matrix factorization (CMF) (Singh & Gordon, 2008) simultaneously factorizes both rating matrix and content matrix with shared item latent factors. SVDFeature (Chen et al., 2012) combines content features with collaborative filtering. The latent factors are integrated with user, item, and global features.

In Du et al. (2014), topic modeling (Blei, Ng, & Jordan, 2003) is utilized to learn topics of users based on the content of their attended events, and then the similarity between topic factor of user and events is calculated, which is an important component of their method. Recently, Zhang and Wang (2015) explicitly addressed the cold-start problem in event recommendation by modeling the event content text. Liao and Chang (2016) proposed a rough set based association rule approach. Sun, Wang, Cheng, and Fu (2015) integrated sentiment information from affective texts into recommendation models. The cold-start problem in tag recommendation is studied in Martins, Belém, Almeida, and Gonçalves (2016). Our work is different from the existing work in that we do not use any textual content information to tackle the cold-start problem. Instead, we formally model group information from the user and event perspectives and demonstrate the advantage of utilizing the dual-view of groups in the cold-start settings.

3. Data analysis

In this section, we provide an analysis of the real-world datasets that we collected from *Meetup* for four cities in U.S.: New York, San Francisco, Washington DC and Chicago, which are among the most active cities for the *Meetup* community.⁴ The detailed statistics and information of the datasets are presented in Section 5.1.

Fig. 1(a) shows the cumulative percentage of users who join a given number of groups, i.e., number of groups (n) vs. the percentage of users joining less or equal than n groups. This information is depicted for all four cities and a consistent

⁴ <http://priceconomics.com/what-meetups-tell-us-about-america/>.

Table 1

Average number of groups per user in four cities.

| New York | San Francisco | Washington DC | Chicago |
|----------|---------------|---------------|---------|
| 9.54 | 9.88 | 8.63 | 7.78 |

Table 2

Notations.

| | |
|--------------------|---|
| m, n, f | Total number of users, events, and latent factors, respectively |
| G_u | The set of groups that user u belongs to |
| (u, i, j) | A preference triplet indicating user u prefers event i over event j |
| D_s | The set that contains all the preference triplets |
| K | The set of (u, i) pairs with known ratings |
| $s_{u, i}$ | Ranking score of event i for user u |
| $x_{u, i, j}$ | Difference of ranking scores between event i and j for user u |
| \mathbf{p}_u | Latent factors for user u |
| \mathbf{q}_i | Latent factors for event i |
| \mathbf{r}_g | User-oriented latent factor for group g |
| \mathbf{t}_g | Event-oriented latent factor for group g that organizes event |
| \mathbf{v}_i | Latent factor for venue that host event i |
| \mathbf{y}_i | Event-oriented latent factor related to the day of the week for event i |
| \mathbf{z}_i | Event-oriented latent factor related to the period of the day for event i |
| c_i | Popularity for event i |
| $d_{u, i}$ | Normalized geo-distance between user u and event i |
| β_c, β_d | Event popularity bias and Geo-distance bias |
| λ, γ | Regularization parameters for latent factors and bias respectively |

pattern emerges from the data. In all four cities, approximately 80% of the users have joined at least one group. Around 5% of the users join one group and a significant majority join between 1 and 30 groups. There are certain users who have joined more than 30 groups, but the percentage of such users is very small. The group membership in the user-base is an important indication that group is an essential feature for event recommendation. Table 1 provides the average number of groups per user in the four cities. We observe from our dataset that in all four cities, there were a couple of users who had joined 200 groups, which is the maximum allowed by *Meetup*. There is also a remarkable consistency in the average number of groups per user in the four cities, which is in the range of 7–10. We also observed that *all* events in our dataset are organized by a group, i.e. the group information is always present in the RSVP. Therefore, a group is a critical contextual parameter that is associated with the majority of users and *all* events.

We also analyzed the data collected for the four cities with respect to relationship between users and venues. As indicated in Fig. 1(b), most users usually attend events at a limited number of venues, which means that there is an implicit relationship between the user and venue. To gain a further insight into the user-venue relationship, we look at a random user in the New York dataset. The user has three RSVPs that are for three different venues. The venues associated with the user's RSVPs are: *230 Fifth* (bar and lounge), *Madison Square Tavern* (restaurant and event center with a full-scale bar) and *Croton Reservoir Tavern* (up-scale restaurant and bar that hosts private parties). The three venues associated with the user are similar based on the fact that all of them are upscale restaurants with bars, and two of them host private events. This observation indicates that the characteristics of the venue may affect the user's RSVP for the event, which motivates us to include the event venue as one of the parameters into our model in Section 4.3.

4. Event recommendation models

As discussed earlier, we tackle the task of event recommendation as the top- N item recommendation by providing a user with a ranked list of events. Our proposed approach is based on the latent factor model with pairwise ranking. In the following subsections, we present our proposed model by considering the dual role of group influence. We then incorporate contextual information into the model, such as venue, event popularity, temporal influence and geographical distance. Table 2 lists the notations used in this paper.

4.1. Pairwise ranking

In EBSNs, most users only respond to a small portion of events since they may only be aware of a few events. In addition, there exist much more positive examples than negative ones. Many users just ignore RSVPs if they are not interested in attending the events. Consequently, there are many unobserved user-event pairs, which are a mixture of real negative feedback (the user is not interested in attending) and missing values (the user might attend if she is aware of the event). Therefore, instead of performing a point-wise RSVP prediction, we use a pairwise ranking approach to learn the preferences of users on events.

Formally, given a user u , if item i is preferred over item j , we have a preference instance $(u, i, j) \in D_S$ where D_S is the whole set of preference instances. In EBSNs, the preference instances can be derived from three types of relations between items given user u : (1) RSVP “Yes” is preferred over RSVP “No”; (2) RSVP “Yes” is preferred over unobserved RSVP; (3) unobserved RSVP of an event organized by the user’s group is preferred over RSVP “No”. Let $s(u, i)$ represent the ranking score of item i for user u and denote:

$$x(u, i, j) = s(u, i) - s(u, j)$$

The pairwise ranking optimization criterion is the log likelihood of the observed preferences, which can then be defined as

$$\max_{\Theta} \mathcal{L}(\Theta) = \sum_{(u,i,j) \in D_S} \log \sigma(x(u, i, j)) - \text{Reg}(\Theta) \quad (1)$$

where $\sigma(x)$ defines the probability of pairwise preference, i.e., the probability of item i being preferred over j given their ranking score difference $x(u, i, j)$. $\sigma(x)$ is a monotonically increasing function with respect to the argument $x(u, i, j)$. The intuitive explanation of Eq. (1) is that if item i is preferred over j for user u , the difference between their ranking scores $s(u, i)$ and $s(u, j)$ is maximized since $\log \sigma(x)$ is a monotonically increasing function. As a result, item i is more preferable than item j . In the above equation, Θ is the set of all model parameters and $\text{Reg}(\Theta)$ is a regularization term to prevent overfitting. In this paper, we use L2 regularization, since the L2-regularization terms are differentiable, allowing us to apply gradient-based optimization methods.

Since $\sigma(x)$ is a probability function while being monotonically increasing, the Logistic function defined as follows is a natural choice:

$$\sigma(x) = \frac{1}{1 + \exp(-x)}$$

In fact, the choice of the Logistic function in Eq. (1) would lead to the widely used Bayesian Personalized Ranking (BPR) optimization criterion (Rendle et al., 2009) in recommender systems. The objective function of BPR is shown as Eq. (3) in Table 3. Theoretically, optimizing for the above BPR is a smoothed version of optimizing for the well-known ranking measure Area under the ROC Curve (AUC) by approximating the non-differentiable Heaviside function by the differentiable Logistic function $\sigma(x)$. See Rendle et al. (2009) for a more detailed explanation. On the other hand, the use of logistic function to model pairwise preference probability is a type of Bradley-Terry models (Agresti & Kateri, 2014) where exponential score functions are used.

Alternatively, we can model the pairwise preference probability $\sigma(x)$ by the Probit function which is a popular specification for an ordinal or a binary response model in Statistics (McCullagh & Nelder, 1989). The Logistic and Probit are both sigmoid functions with a domain between 0 and 1, which makes them both quantile functions - i.e., inverses of the cumulative distribution function (CDF) of a probability distribution. In fact, the Logistic is the quantile function of the Logistic distribution, while the Probit is the quantile function of the Gaussian distribution defined as follows:

$$\sigma(x) = \Phi(x) = \int_{-\infty}^x \mathcal{N}(x) dx = \int_{-\infty}^x \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx \quad (2)$$

where $\Phi(x)$ is the cumulative distribution function of Gaussian distribution. $\mathcal{N}(x)$ is the probability density function of the Gaussian distribution. In this paper, we set $\mu = 0$ and $\sigma^2 = 1$ yielding the standard Gaussian distribution. Both Logistic and Probit functions have a similar ‘S’ shape. The Logistic has a slightly flatter tail while the Probit curve approaches the axes more quickly. In the Probit function, if we increase the variance the curve will become flatter and elongated. The experiments in Section 5 compare the performance of the two variants of the proposed model.

4.2. Group-aware latent factor model

The latent factor model is one of the most successful collaborative filtering models, which jointly maps the users and items into a shared latent space of a much lower dimensionality. Here we use latent factor model to characterize the ranking scores $s(u, i)$ and $s(u, j)$ in Eq. (1). Formally, users and events are projected into a shared f -dimensional latent space, where $f \ll \min(m, n)$: m is the number of users and n is the number of events. In the most basic form, user u is mapped to a latent factor vector $\mathbf{p}_u \in \mathbb{R}^f$, and event i is mapped to a latent factor vector $\mathbf{q}_i \in \mathbb{R}^f$. The inner product of \mathbf{p}_u and \mathbf{q}_i is used to compute the predicted ranking score of user u on event i such as $s_{u,i} = \mathbf{p}_u^T \mathbf{q}_i$. Similarly, we have $s_{u,j} = \mathbf{p}_u^T \mathbf{q}_j$ for event j where \mathbf{q}_j is the latent factor for j .

Based on our data analysis in Section 3, a large majority of users are associated with at least one group and each event is organized by one group. These observations suggest that considering the group influence may improve the accuracy of event recommendation. We can view the group influence from a dual perspective: user-oriented and event-oriented. The user-oriented perspective regards a group as a topic of interest so that users associated with a group are interested in the same topic with the group. On the other hand, the event-oriented perspective views a group as an organizing entity. The events organized by the same group have the same organizing style such as logistics, event planning, structure, quality of talks, etc. These two perspectives complement each other and together they form a complete view of a group.

Table 3Objective functions $\mathcal{L}(\Theta)$ for BPR, GLFM, GLFM-V, GLFM-VPD, and GLFM-VPDT, respectively.

$$\max_{P,Q} \sum_{(u,i,j) \in D_s} \log \sigma(x(u, i, j)) - \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 \right) \quad (3)$$

$$\max_{P,Q,R,T} \sum_{(u,i,j) \in D_s} \log \sigma(x(u, i, j)) - \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 + \sum_g \|\mathbf{r}_g\|_2^2 + \sum_i \|\mathbf{t}_g\|_2^2 \right) \quad (4)$$

$$\max_{P,Q,R,T} \sum_{(u,i,j) \in D_s} \log \sigma(x(u, i, j)) - \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 + \sum_g \|\mathbf{r}_g\|_2^2 + \sum_i \|\mathbf{t}_g\|_2^2 + \sum_i \|\mathbf{v}_i\|_2^2 \right) \quad (5)$$

$$\max_{P,Q,R,T} \sum_{(u,i,j) \in D_s} \log \sigma(x(u, i, j)) - \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 + \sum_g \|\mathbf{r}_g\|_2^2 + \sum_i \|\mathbf{t}_g\|_2^2 + \sum_i \|\mathbf{v}_i\|_2^2 \right) - \gamma (\beta_c^2 + \beta_d^2) \quad (6)$$

$$\max_{P,Q,R,T,Y,Z} \sum_{(u,i,j) \in D_s} \log \sigma(x(u, i, j)) - \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 + \sum_g \|\mathbf{r}_g\|_2^2 + \sum_i \|\mathbf{t}_g\|_2^2 + \sum_i \|\mathbf{v}_i\|_2^2 + \sum_i \|\mathbf{y}_i\|_2^2 + \sum_i \|\mathbf{z}_i\|_2^2 \right) - \gamma (\beta_c^2 + \beta_d^2) \quad (7)$$

We propose the group-aware latent factor model (GLFM) to model user preference by encoding the dual perspective of group influence. Mathematically, for group g , we use \mathbf{r}_g and \mathbf{t}_g to model its user-oriented and event-oriented characteristics, respectively. Since a user could be a member of multiple groups, we average all the user-oriented latent vectors \mathbf{r}_g of these groups and use it to influence the user latent factor. Similarly, we use the event-oriented latent factor \mathbf{t}_g of the group that organizes the event to influence the event latent factor. Let G_u be the set of groups that user u belongs to. Let $g \in G_u$ be a specific group that includes user u . Let \mathbf{t}_g denote the latent factor of the group that organizes event i . Incorporated with influence from groups, the predicted ranking score for event i given user u is now computed with both \mathbf{r}_g and \mathbf{t}_g , shown as follows.

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g) \quad (8)$$

The ranking score $s_{u,j}$ for event j given user u can be similarly calculated. The objective function is shown as Eq. (4) in Table 3.

It is worth noting that by considering the group information, GLFM addresses the cold-start problems for both new events and new users that do not appear in training data. When a new event i is just released in an EBSN, we do not have any information about \mathbf{q}_i , but the event-oriented group latent factor \mathbf{t}_g is not empty since we know which group organizes the event. The latent factor of the group is learned from the training data. Intuitively, if the group has an excellent track record of organizing events like having great talks and good event planning, users may prefer the events organized by this group. Similarly, when a new user u has not responded to any RSVPs, we do not have the latent factor \mathbf{p}_u , but we may know what groups she is associated with and thus can use $\frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g$ for prediction/ranking. These latent factors capture the user-oriented characteristics of the groups such as topics of interest. We can view $\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g$ as a kind of smoothed version of user latent factor smoothed by the groups that the user belongs to. The group influence serves as the background model and is crucial when \mathbf{p}_u is empty. Similarly, $\mathbf{q}_i + \mathbf{t}_g$ can be viewed as the smoothed version of event latent factor. With these latent group factors, we can tackle new users and new items.

4.3. Event venue

Each event is held at a venue. Some groups often organize events at the same or a similar venue, indicating a correlation between the event group and the event venue. The event venue may affect the attendance of the event. For example, some venues have a great facility, or they are at a convenient location, which can attract more people in general. Some venues can accommodate special needs of certain users such as being pets or kids friendly. Some venues are specialized for certain types of events such as ballroom dance or tennis games.

We introduce venue latent factors to exploit event venues for more accurate event recommendation. We treat the venue as an attribute of the event and augment our model with a latent factor \mathbf{v}_i for the venue that hosts event i . The model that includes the influence of venue is GLFM-V. By incorporating the venue influence, the predicted ranking score of event

i given user u is now defined as

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g + \mathbf{v}_i) \quad (9)$$

The objective function is shown as Eq. (5) in Table 3.

4.4. Event popularity and geographical distance

In EBSNs, some events have general themes such as entrepreneurship, while others have niche topics such as *Minecraft*. A hypothesis is users may be more likely to RSVP on popular/mainstream events than on unpopular/niche events. We measure the event popularity by the number of people who RSVP for the event. An event that has a higher number of RSVPs is considered to be more popular. We include event popularity as a ranking bias and perform feature scaling while considering this in conjunction with the other features. By incorporating the popularity bias, the predicted ranking score of event i given user u is as follows:

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g + \mathbf{v}_i) + \beta_c c_i \quad (10)$$

where c_i is the popularity bias for event i and β_c is the weight of the bias which is learned from the training data.

Geographical distance is another important consideration while recommending products and services that require the user to travel to the location. We incorporate the geographical distance in our model by computing the Haversine distance (Shumaker & Sinnott, 1984) from the user latitude-longitude and venue latitude-longitude data. We calculate the logarithm of this distance and model it as a ranking bias.

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g + \mathbf{v}_i) + \beta_c c_i + \beta_u d_{ui} \quad (11)$$

where $d_{u,i}$ is the normalized logarithm geo-distance between user u and venue that hosts event i , and β_u is the geo-distance bias parameter associated with the user that is learned from the training data. The objective function is provided in Eq. (6) in Table 3. In the equation, γ is the regularization parameter used to prevent overfitting. The model that augments the group and venue latent factors with the event popularity and geo-distance bias is called GLFM-VPD.

4.5. Temporal influence

Events are organized during certain days of the week and at certain times of the day. Some events are organized in the day between 9am and 5pm, whereas others are organized in the evening after 5pm, so people can attend after work. Events that are targeted towards working individuals are generally organized on the weekends. We incorporate the temporal influence into our model using two types of latent time factors: one for the day of the week, \mathbf{y}_i , which is associated with the event i , and another for the period of the day, \mathbf{z}_i , for event i . The day of the week is derived from weekday that the event is scheduled on, whereas the period of the day is mapped to one of two time-slots: “Day” if the event time is between 9am and 5pm, and “Evening” if the event is scheduled after 5pm.

With the inclusion of the temporal influence parameters, the predicted ranking score of event i given user u is now defined as

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g + \mathbf{v}_i + \mathbf{y}_i + \mathbf{z}_i) + \beta_c c_i + \beta_u d_{ui} \quad (12)$$

where the \mathbf{y}_i parameter that models the influence of the day of the week, and the \mathbf{z}_i models the influence of the period of the day. The objective function is provided in Eq. (7) in Table 3 with the model denoted by GLFM-VPDT.

4.6. Parameter estimation

We estimate the model parameters of the proposed models by using stochastic gradient descent (SGD) algorithm (Bottou, 2010). In this case, an update is performed for each preference instance $(u, i, j) \in D_S$. Since we deal with maximization problems, the parameters are learned by moving in the direction of the gradient with a learning rate α in an iterative manner as follows.

$$\Theta \leftarrow \Theta - \alpha \frac{\partial \mathcal{L}}{\partial \Theta} \quad (13)$$

Table 4
Stochastic gradient descent updates for GLFM-VPDT.

$$\begin{aligned}
\mathbf{p}_u &\leftarrow \mathbf{p}_u + \alpha \cdot \left(\omega_{u,i,j} \cdot ((\mathbf{q}_i + \mathbf{t}_{g(i)} + \mathbf{v}_i + \mathbf{y}_i + \mathbf{z}_i) - (\mathbf{q}_j + \mathbf{t}_{g(j)} + \mathbf{v}_j + \mathbf{y}_j + \mathbf{z}_j)) - \lambda \cdot \mathbf{p}_u \right) \\
\forall g \in G_u : \mathbf{r}_g &\leftarrow \mathbf{r}_g + \alpha \cdot \left(\omega_{u,i,j} \cdot ((\mathbf{q}_i + \mathbf{t}_{g(i)} + \mathbf{v}_i + \mathbf{y}_i + \mathbf{z}_i) - (\mathbf{q}_j + \mathbf{t}_{g(j)} + \mathbf{v}_j + \mathbf{y}_j + \mathbf{z}_j)) - \lambda \cdot \mathbf{r}_g \right) \\
\mathbf{q}_i &\leftarrow \mathbf{q}_i + \alpha \cdot (\omega_{u,i,j} \cdot (\mathbf{p}_u + \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{q}_i) \\
\mathbf{q}_j &\leftarrow \mathbf{q}_j + \alpha \cdot (\omega_{u,i,j} \cdot (-\mathbf{p}_u - \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{q}_j) \\
\mathbf{t}_{g(i)} &\leftarrow \mathbf{t}_{g(i)} + \alpha \cdot (\omega_{u,i,j} \cdot (\mathbf{p}_u + \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{t}_{g(i)}) \\
\mathbf{t}_{g(j)} &\leftarrow \mathbf{t}_{g(j)} + \alpha \cdot (\omega_{u,i,j} \cdot (-\mathbf{p}_u - \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{t}_{g(j)}) \\
\mathbf{v}_i &\leftarrow \mathbf{v}_i + \alpha \cdot (\omega_{u,i,j} \cdot (\mathbf{p}_u + \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{v}_i) \\
\mathbf{v}_j &\leftarrow \mathbf{v}_j + \alpha \cdot (\omega_{u,i,j} \cdot (-\mathbf{p}_u - \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{v}_j) \\
\mathbf{y}_i &\leftarrow \mathbf{y}_i + \alpha \cdot (\omega_{u,i,j} \cdot (\mathbf{p}_u + \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{y}_i) \\
\mathbf{y}_j &\leftarrow \mathbf{y}_j + \alpha \cdot (\omega_{u,i,j} \cdot (-\mathbf{p}_u - \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{y}_j) \\
\mathbf{z}_i &\leftarrow \mathbf{z}_i + \alpha \cdot (\omega_{u,i,j} \cdot (\mathbf{p}_u + \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{z}_i) \\
\mathbf{z}_j &\leftarrow \mathbf{z}_j + \alpha \cdot (\omega_{u,i,j} \cdot (-\mathbf{p}_u - \sum_{ug} \mathbf{r}_g) - \lambda \cdot \mathbf{z}_j) \\
\text{for } i, \beta_c &\leftarrow \beta_c + \alpha \cdot (\omega_{u,i,j} \cdot (c_i + c_j) - \gamma \cdot \beta_c) \\
\beta_d &\leftarrow \beta_d + \alpha \cdot (\omega_{u,i,j} \cdot (d_{ui} + d_{uj}) - \gamma \cdot \beta_d) \\
\text{for } j, \beta_c &\leftarrow \beta_c + \alpha \cdot (\omega_{u,i,j} \cdot (-c_i - c_j) - \gamma \cdot \beta_c) \\
\beta_d &\leftarrow \beta_d + \alpha \cdot (\omega_{u,i,j} \cdot (-d_{ui} - d_{uj}) - \gamma \cdot \beta_d)
\end{aligned}$$

By plugging our pairwise ranking optimization criterion in Eq. (1) into Eq. (13), we obtain

$$\Theta \leftarrow \Theta - \alpha \left(\frac{1}{\sigma(x(u, i, j))} \frac{\partial \sigma(x(u, i, j))}{\partial \Theta} - \frac{\partial \text{Reg}(\Theta)}{\partial \Theta} \right) \quad (14)$$

The algorithm repeatedly iterates over the training data and updates the model parameters in each iteration until convergence. Based on Eq. (14), we can derive the SGD updates for GLFM-VPDT as shown in Table 4. The updates for other proposed latent factor models (e.g., GLFM, GLFM-V, GLFM-VPD) are similar. In the table, $\omega_{u,i,j}$ is defined in order to simplify the notation. For the model based on the Logistic function,

$$\omega_{u,i,j} = \frac{e^{-x(u,i,j)}}{1 + e^{-x(u,i,j)}}$$

For the model based on the Probit function:

$$\omega_{u,i,j} = \frac{\mathcal{N}(x(u, i, j))}{\Phi(x(u, i, j))}$$

where $\mathcal{N}(\cdot)$ and $\Phi(\cdot)$ are defined in Eq. (2).

As introduced in Section 4.1, the preference instances can be derived from three types of relations between items given a user based on RSVP “Yes”, RSVP “No”, and missing RSVP. Thus, we generate the preference instances (u, i, j) from the training data based on the following strategy:

- If the user has any positive RSVPs, we randomly sample a RSVP “Yes” and then randomly sample a RSVP “No” from the same user to form the preference pair.
- If there is no negative RSVP for the user, we randomly sample a missing RSVP from the user. This pairing is based on the assumption that a RSVP with unknown preference is negative when paired with a true positive example.
- If the user has no positive RSVP, we pair a random unknown RSVP of an event organized by one of the user’s groups with a random negative RSVP from the same user. This pairing is based on the assumption that an unknown preference for a RSVP of an event organized by a group that the user belongs to, is preferred over a true negative example.

Section 5.3.3 investigates an alternative preference generation strategy without assuming that the unobserved RSVPs are preferred over the observed RSVP “No”. Once a sufficient number of preference instances are sampled, we randomly shuffle them to avoid bias for certain users. The model is then trained on these permuted instances by SGD. The learned model parameters are then applied to the test users for the top- N event recommendation based on descending order of the ranking score $S_{u,i}$.

5. Experiments

We evaluate the proposed dual perspective group-aware model and its variants on real-world datasets collected from *Meetup*. We compare the results of our models against the state-of-the-art recommendation techniques. In addition to performing a comparison in regular settings, we also compare our models with the baseline methods in cold-start scenarios. We present the results in this section and discuss our findings in detail.

Table 5
Data statistics for four cities.

| City | RSVPs | Sparsity | Positive | Negative | Users | Events | Groups | Venues |
|---------------|--------|----------|----------|----------|-------|--------|--------|--------|
| New York | 50,150 | 0.9989 | 49,163 | 987 | 1397 | 35,179 | 1326 | 1696 |
| San Francisco | 24,923 | 0.9984 | 23,848 | 1075 | 1147 | 13,938 | 748 | 1075 |
| DC | 23,688 | 0.9968 | 23,205 | 483 | 635 | 11,906 | 503 | 845 |
| Chicago | 12,598 | 0.9976 | 11,782 | 816 | 599 | 8819 | 433 | 853 |

5.1. Data collection

As introduced in Section 3, we collected RSVP data from *Meetup* for events organized in four cities in the U.S.: New York, San Francisco, Washington DC and Chicago. The RSVP data was collected by using the Meetup API⁵ between the time periods January 2016 and May 2016. We filtered the dataset for each city to retain only RSVPs associated with users having greater than 5 RSVPs. The statistics of the data are given in Table 5. These four cities represent different geographic regions of the U.S. and they have varied statistics as shown in the table.

Table 5 also provides statistics of the RSVPs, including the breakup of the RSVPs into the positive and negative ones. We observe that for all four cities, the positive RSVPs far exceed the negative ones. This indicates that users generally respond when they are interested in attending an event. Users who intend to respond with a negative or *no* RSVP for an event generally ignore the event and do not provide a response. This observation justifies our pairwise learning approach, which utilizes both negative and unobserved RSVPs by forming preference pairs instead of performing pointwise prediction. Section 5.3 includes the comparison of the experimental results of different methods.

5.2. Experimental setup

We sorted the data in chronological order of event time so that we can train the model on past events and recommend future ones. The sorted datasets are then split to use 80% as the training set and 20% as the test set for each city. We apply the sampling strategy introduced in Section 4.6 to generate preference instances for model training. The learned model parameters are applied to the test users to generate a ranking score for the events for each user based on $s_{u, i}$. The evaluation metrics include $P@5$, $P@10$, $R@5$, $R@10$, $NDCG@5$, $NDCG@10$, and $MAP@10$ (Manning, Raghavan, Schütze et al., 2008). These are common metrics for top- N recommendations.

We compare the proposed models with the following baseline methods. We use *Librec*.⁶ a widely used recommender library, to obtain results for the baseline methods. We set both the regularization parameters λ and γ to 0.025 and the learning rate in SGD is $\alpha = 0.05$. We use the same parameter values with the existing methods (to the extent possible) as we have used with the proposed methods.

- *Item Mean*: The ranking score of an event is predicted on the basis of the mean of the event RSVPs in the training set.
- *User KNN* (Breese, Heckerman, & Kadie, 1998): User-based K -Nearest Neighborhood collaborative filtering method that predicts the user preference based on the similarity with the K nearest users calculated using Pearson's correlation.
- *Item KNN* (Sarwar, Karypis, Konstan, & Riedl, 2001): Item-based K -Nearest Neighborhood collaborative filtering method that predicts the user preference based on the similarity with the K nearest items calculated using Pearson's correlation.
- *Group-Membership*: This is a naive method that utilizes the user group membership data to recommend events organized by the groups that the users belong to. If the user does not belong to any group, then this model recommends the most popular events (based on the RSVP count) to the user. The user group membership data is obtained using the Meetup API.
- *Biased-MF* (Koren et al., 2009): Basic matrix factorization that includes global mean, user bias and event bias.
- *BPR-MF* (Rendle et al., 2009): Bayesian Personalized Ranking method that utilizes pairwise loss to provide top- N item recommendation using matrix factorization (MF).
- *SVD++* (Koren, 2008): A state-of-the-art matrix factorization method that incorporates implicit feedback from the user for a superior accuracy.
- *SVDFeature*.⁷ State-of-the-art model that incorporates domain-specific features to SVD++. We configure this toolkit to utilize the group and venue information from the dataset. We indicate the group as both a user and event feature, and the venue as only an event feature.

We evaluate the following proposed models by integrating influences from multiple factors: group-aware (G), venue influence (V), popularity influence (P), distance influence (D) and temporal influence (T). Our models are also varied with the choice of the pairwise probability functions: Logistic (Logit) or Probit.

⁵ http://www.meetup.com/meetup_api/.

⁶ <http://www.librec.net>.

⁷ <http://svdfeature.apexlab.org/>.

- *LFM-V-Logit*: This model incorporates the influence of only the event venue (V) into the BPR with matrix factorization (BPR-MF). The logistic function is used for the pairwise ranking.
- *LFM-T-Logit*: This model incorporates the temporal influence into the BPR with matrix factorization (BPR-MF), which utilizes the logistic function.
- *LFM-P-Logit*: This model incorporates the influence of the event popularity (P) into the BPR with matrix factorization (BPR-MF), which utilizes the logistic function.
- *LFM-D-Logit*: This model incorporates the influence of the geographical distance (D) into the BPR with matrix factorization (BPR-MF), which utilizes the logistic function.
- *GLFM-Logit*: This model considers the influence of the dual perspective of groups (G) and uses the logistic function for the pairwise ranking.
- *GLFM-V-Logit*: This model considers the dual-perspective groups (G) and the influence of event venue (V) with the logistic function.
- *GLFM-VPD-Logit*: This model includes the dual perspective of groups (G), event venue (V), event popularity (P) and geographical distance (D) with the logistic pairwise function.
- *GLFM-VPDT-Logit*: This model includes all the information – dual perspective of groups (G), event venue (V), event popularity (P), geographical distance (D), and temporal influence (T) – with the logistic pairwise function.
- *LFM-V-Probit*: This model is similar to *LFM-V-Logit* but with the Probit pairwise probability function.
- *LFM-T-Probit*: This model is similar to *LFM-T-Logit* but with the Probit pairwise probability function.
- *LFM-P-Probit*: This model is similar to *LFM-P-Logit* but with the Probit pairwise probability function.
- *LFM-D-Probit*: This model is similar to *LFM-D-Logit* but with the Probit pairwise probability function.
- *GLFM-Probit*: This model is similar to *GLFM-Logit* but with the Probit pairwise probability function.
- *GLFM-V-Probit*: This model is similar to *GLFM-V-Logit* but with the Probit pairwise probability function.
- *GLFM-VPD-Probit*: This model is similar to *GLFM-VPD-Logit* but with the Probit pairwise probability function.
- *GLFM-VPDT-Probit*: This model is similar to *GLFM-VPDT-Logit* but with the Probit pairwise probability function.
- *GLFM-VPDT-Pointwise*: This model includes the dual perspective of group (G), event venue (V), event popularity (P), geographical distance (D), and temporal influence (T), utilizing a point-wise loss function. The ranking score for the user u on the event i is as follows. The objective function considers the actual RSVP of the user $a_{u,i}$, with the value 1 if the user provided an affirmative RSVP, and 0 for a negative or missing RSVP.

$$s_{u,i} = \left(\mathbf{p}_u + \frac{1}{|G_u|} \sum_{g \in G_u} \mathbf{r}_g \right)^T (\mathbf{q}_i + \mathbf{t}_g + \mathbf{v}_i + \mathbf{y}_i + \mathbf{z}_i) + \beta_c c_i + \beta_u d_{ui} \quad (15)$$

$$\min_{P,Q,R,T,Y,Z} \sum_{(u,i) \in K} (s_{u,i} - a_{u,i})^2 + \lambda \left(\sum_u \|\mathbf{p}_u\|_2^2 + \sum_i \|\mathbf{q}_i\|_2^2 + \sum_g \|\mathbf{r}_g\|_2^2 + \sum_i \|\mathbf{t}_g\|_2^2 + \sum_i \|\mathbf{v}_i\|_2^2 + \sum_i \|\mathbf{y}_i\|_2^2 + \sum_i \|\mathbf{z}_i\|_2^2 \right) + \gamma (\beta_c^2 + \beta_d^2) \quad (16)$$

5.3. Results

In the following subsections, we first compare the proposed models with the baseline methods. The effect of dimensionality of latent factor space is also investigated. We then conduct experiments in cold-start settings to evaluate the proposed methods.

5.3.1. Baseline comparisons

The results of baseline comparisons are presented in [Tables 6–9](#) for the four cities, respectively. The best results in each evaluation metric are highlighted in boldface. From the tables, we can make the following observations.

- The proposed dual-perspective group-aware models (i.e., *GLFM-Logit*, *GLFM-Probit*, *GLFM-V-Logit*, *GLFM-V-Probit*, *GLFM-VPD-Logit*, *GLFM-VPD-Probit*, *GLFM-VPDT-Logit*, *GLFM-VPDT-Probit*) substantially outperform the methods that do not consider group information. The best results in all the four cities are achieved by the proposed latent factor models with a large margin of improvement.
- Excluding the proposed group-aware models, the venue-aware latent factor models (*LFM-V*) perform the best. These results validate our observation in [Section 3](#) that there may exist a correlation between the user and venue. Event venue is an important consideration while deciding to attend an event, as groups generally host events at the same or similar venues.
- The latent factor models that only consider the temporal influence, event popularity and geographical distance (*LFM-T*, *LFM-P*, *LFM-D*) yield mediocre results for all the four cities. Out of the three models, the model with the temporal influence yields the best results, followed by the model that considered the event popularity and geographical distance. The

Table 6
Experimental results of baseline comparisons for New York.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|---------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Item Mean | 0.0062 | 0.0033 | 0.0158 | 0.0164 | 0.0142 | 0.0163 | 0.0138 |
| UserKNN | 0.0187 | 0.0102 | 0.0403 | 0.0428 | 0.0372 | 0.0419 | 0.0336 |
| ItemKNN | 0.0445 | 0.0257 | 0.0731 | 0.0785 | 0.0679 | 0.0728 | 0.0558 |
| Biased-MF | 0.0002 | 0.0002 | 0.0009 | 0.0017 | 0.0007 | 0.0008 | 0.0004 |
| BPR-MF | 0.2340 | 0.1743 | 0.2559 | 0.3242 | 0.2819 | 0.3004 | 0.2594 |
| SVD++ | 0.0003 | 0.0003 | 0.0015 | 0.0024 | 0.0011 | 0.0012 | 0.0008 |
| SVDFeature | 0.1606 | 0.1597 | 0.1877 | 0.2030 | 0.2223 | 0.2478 | 0.1505 |
| Group-Membership | 0.4063 | 0.3806 | 0.2155 | 0.3581 | 0.3791 | 0.4247 | 0.4029 |
| GLFM-VPDT-Pointwise | 0.0036 | 0.0030 | 0.0022 | 0.0050 | 0.0011 | 0.0013 | 0.0023 |
| LFM-V-Logit | 0.4491 | 0.4456 | 0.16641 | 0.3170 | 0.2580 | 0.2849 | 0.4481 |
| LFM-T-Logit | 0.2540 | 0.2177 | 0.1805 | 0.2130 | 0.2565 | 0.2682 | 0.2253 |
| LFM-P-Logit | 0.2020 | 0.1865 | 0.1437 | 0.1943 | 0.2005 | 0.2116 | 0.1919 |
| LFM-D-Logit | 0.1870 | 0.1634 | 0.1582 | 0.1698 | 0.1905 | 0.2145 | 0.1554 |
| GLFM-Logit | 0.7153 | 0.6623 | 0.3463 | 0.5615 | 0.4284 | 0.4731 | 0.7107 |
| GLFM-V-Logit | 0.7180 | 0.6760 | 0.3327 | 0.5628 | 0.4272 | 0.4717 | 0.7143 |
| GLFM-VPD-Logit | 0.7193 | 0.6707 | 0.3242 | 0.5413 | 0.4277 | 0.4724 | 0.7135 |
| GLFM-VPDT-Logit | 0.7177 | 0.6777 | 0.3292 | 0.5631 | 0.4218 | 0.4658 | 0.7124 |
| LFM-V-Probit | 0.4284 | 0.4202 | 0.1687 | 0.3070 | 0.2478 | 0.2736 | 0.4271 |
| LFM-T-Probit | 0.2446 | 0.2357 | 0.1881 | 0.1917 | 0.1938 | 0.2133 | 0.2349 |
| LFM-P-Probit | 0.2166 | 0.2070 | 0.1549 | 0.1706 | 0.1725 | 0.1816 | 0.2054 |
| LFM-D-Probit | 0.1871 | 0.1760 | 0.1410 | 0.1476 | 0.1522 | 0.1643 | 0.1606 |
| GLFM-Probit | 0.6886 | 0.6555 | 0.3074 | 0.5282 | 0.4045 | 0.4467 | 0.6854 |
| GLFM-V-Probit | 0.7217 | 0.6794 | 0.3364 | 0.5749 | 0.4248 | 0.4691 | 0.7172 |
| GLFM-VPD-Probit | 0.7257 | 0.6856 | 0.3287 | 0.5767 | 0.4268 | 0.4713 | 0.7207 |
| GLFM-VPDT-Probit | 0.7397 | 0.6908 | 0.3353 | 0.5589 | 0.4418 | 0.4879 | 0.7353 |

Table 7
Experimental results of baseline comparisons for San Francisco.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|---------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Item Mean | 0.0055 | 0.0031 | 0.0108 | 0.0118 | 0.0097 | 0.0036 | 0.0021 |
| UserKNN | 0.0321 | 0.0253 | 0.0732 | 0.1165 | 0.0592 | 0.0704 | 0.0486 |
| ItemKNN | 0.0548 | 0.0370 | 0.1267 | 0.1776 | 0.1219 | 0.1310 | 0.0950 |
| Biased-MF | 0.0004 | 0.0003 | 0.0010 | 0.0024 | 0.0004 | 0.0011 | 0.0006 |
| BPR-MF | 0.2124 | 0.1700 | 0.2552 | 0.3661 | 0.2991 | 0.3097 | 0.2588 |
| SVD++ | 0.0003 | 0.0003 | 0.0013 | 0.0027 | 0.0007 | 0.0008 | 0.0004 |
| SVDFeature | 0.1499 | 0.1343 | 0.1926 | 0.2180 | 0.1903 | 0.1979 | 0.1366 |
| Group-Membership | 0.3019 | 0.2746 | 0.2068 | 0.3317 | 0.4063 | 0.4513 | 0.2987 |
| GLFM-VPDT-Pointwise | 0.0043 | 0.0034 | 0.0151 | 0.0236 | 0.0019 | 0.0021 | 0.0041 |
| LFM-V-Logit | 0.3295 | 0.2143 | 0.1848 | 0.2591 | 0.2208 | 0.2412 | 0.4364 |
| LFM-T-Logit | 0.2277 | 0.2001 | 0.2385 | 0.2405 | 0.2164 | 0.2182 | 0.2266 |
| LFM-P-Logit | 0.1973 | 0.1821 | 0.1547 | 0.1714 | 0.2412 | 0.2589 | 0.1902 |
| LFM-D-Logit | 0.1888 | 0.1709 | 0.1787 | 0.2088 | 0.1612 | 0.1842 | 0.1793 |
| GLFM-Logit | 0.5843 | 0.5110 | 0.3657 | 0.5573 | 0.3529 | 0.3897 | 0.5745 |
| GLFM-V-Logit | 0.6355 | 0.5542 | 0.3721 | 0.5506 | 0.3861 | 0.4264 | 0.6248 |
| GLFM-VPD-Logit | 0.6403 | 0.5629 | 0.3697 | 0.5545 | 0.3890 | 0.4296 | 0.6317 |
| GLFM-VPDT-Logit | 0.6186 | 0.5457 | 0.3402 | 0.5085 | 0.3774 | 0.4168 | 0.6108 |
| LFM-V-Probit | 0.3757 | 0.3372 | 0.1914 | 0.3549 | 0.2137 | 0.2361 | 0.3764 |
| LFM-T-Probit | 0.2390 | 0.2280 | 0.2090 | 0.2254 | 0.2272 | 0.2490 | 0.2307 |
| LFM-P-Probit | 0.2134 | 0.1933 | 0.1689 | 0.1815 | 0.2017 | 0.2288 | 0.2033 |
| LFM-D-Probit | 0.1799 | 0.1633 | 0.1500 | 0.1645 | 0.1911 | 0.1939 | 0.1771 |
| GLFM-Probit | 0.6112 | 0.5392 | 0.3666 | 0.5539 | 0.3651 | 0.4032 | 0.5985 |
| GLFM-V-Probit | 0.6026 | 0.5353 | 0.3573 | 0.5444 | 0.3637 | 0.4017 | 0.5948 |
| GLFM-VPD-Probit | 0.6026 | 0.5284 | 0.3612 | 0.5414 | 0.3646 | 0.4026 | 0.5935 |
| GLFM-VPDT-Probit | 0.5826 | 0.5366 | 0.2989 | 0.4848 | 0.3449 | 0.3808 | 0.5772 |

geographical distance has the least impact due to the fact that the RSVPs were considered local to a city, so the distance is not a major consideration while deciding to attend an event. In conjunction with group and venue information, the event popularity, temporal influence and geographical distance may slightly improve the performance, as shown on the New York and San Francisco datasets.

Table 8
Experimental results of baseline comparisons for Washington DC.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|---------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Item Mean | 0.0029 | 0.0021 | 0.0053 | 0.0112 | 0.0058 | 0.0061 | 0.0034 |
| UserKNN | 0.0043 | 0.0027 | 0.0078 | 0.0132 | 0.0068 | 0.0074 | 0.0050 |
| ItemKNN | 0.0312 | 0.0161 | 0.0731 | 0.0773 | 0.0661 | 0.0698 | 0.0566 |
| Biased-MF | 0.0013 | 0.0012 | 0.0039 | 0.0072 | 0.0021 | 0.0023 | 0.0009 |
| BPR-MF | 0.3322 | 0.2633 | 0.3692 | 0.4843 | 0.3915 | 0.4270 | 0.3664 |
| SVD++ | 0.0014 | 0.0019 | 0.0043 | 0.0086 | 0.0030 | 0.0031 | 0.0016 |
| SVDFeature | 0.2281 | 0.2104 | 0.2332 | 0.2556 | 0.2789 | 0.2831 | 0.2175 |
| Group-Membership | 0.3120 | 0.2690 | 0.2482 | 0.3656 | 0.3878 | 0.4321 | 0.3110 |
| GLFM-VPDT-Pointwise | 0.0043 | 0.0040 | 0.0052 | 0.0069 | 0.0020 | 0.0023 | 0.0047 |
| LFM-V-Logit | 0.2880 | 0.2538 | 0.1425 | 0.2099 | 0.1929 | 0.2131 | 0.2959 |
| LFM-T-Logit | 0.2243 | 0.2112 | 0.1629 | 0.1641 | 0.2458 | 0.2506 | 0.2146 |
| LFM-P-Logit | 0.1861 | 0.1854 | 0.1336 | 0.1572 | 0.1974 | 0.2114 | 0.1860 |
| LFM-D-Logit | 0.1402 | 0.1356 | 0.1222 | 0.1279 | 0.1833 | 0.1914 | 0.1390 |
| GLFM-Logit | 0.6909 | 0.5825 | 0.4766 | 0.6605 | 0.4414 | 0.4875 | 0.6820 |
| GLFM-V-Logit | 0.7410 | 0.6167 | 0.5111 | 0.7003 | 0.4753 | 0.5249 | 0.7315 |
| GLFM-VPD-Logit | 0.7301 | 0.6123 | 0.5026 | 0.6873 | 0.4631 | 0.5031 | 0.7100 |
| GLFM-VPDT-Logit | 0.7032 | 0.6153 | 0.4603 | 0.6748 | 0.4373 | 0.4829 | 0.6950 |
| LFM-V-Probit | 0.3469 | 0.3134 | 0.1741 | 0.2632 | 0.2214 | 0.2445 | 0.3514 |
| LFM-T-Probit | 0.2061 | 0.2043 | 0.1956 | 0.2183 | 0.2611 | 0.2675 | 0.2060 |
| LFM-P-Probit | 0.1750 | 0.1646 | 0.1342 | 0.1491 | 0.2361 | 0.2399 | 0.1677 |
| LFM-D-Probit | 0.1651 | 0.1642 | 0.1342 | 0.1507 | 0.2112 | 0.2349 | 0.1640 |
| GLFM-Probit | 0.6756 | 0.5832 | 0.4540 | 0.6580 | 0.4262 | 0.4707 | 0.6690 |
| GLFM-V-Probit | 0.6865 | 0.6080 | 0.4407 | 0.6469 | 0.4357 | 0.4811 | 0.6848 |
| GLFM-VPD-Probit | 0.6516 | 0.5720 | 0.4213 | 0.6026 | 0.4274 | 0.4719 | 0.6587 |
| GLFM-VPDT-Probit | 0.6843 | 0.6043 | 0.4367 | 0.6634 | 0.4181 | 0.4617 | 0.6744 |

Table 9
Experimental results of baseline comparisons for Chicago.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|---------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Item Mean | 0.0005 | 0.0002 | 0.0002 | 0.0002 | 0.0002 | 0.0002 | 0.0001 |
| UserKNN | 0.0252 | 0.0171 | 0.0583 | 0.0757 | 0.0592 | 0.0623 | 0.0476 |
| ItemKNN | 0.0516 | 0.0306 | 0.1166 | 0.1348 | 0.1025 | 0.1150 | 0.0870 |
| Biased-MF | 0.0012 | 0.0010 | 0.0036 | 0.0066 | 0.0034 | 0.0033 | 0.0019 |
| BPR-MF | 0.1653 | 0.1178 | 0.3011 | 0.3867 | 0.2917 | 0.3167 | 0.2672 |
| SVD++ | 0.0015 | 0.0014 | 0.0047 | 0.0087 | 0.0038 | 0.0043 | 0.0023 |
| SVDFeature | 0.1441 | 0.1402 | 0.2754 | 0.2819 | 0.3043 | 0.3109 | 0.1414 |
| Group-Membership | 0.4189 | 0.3536 | 0.3322 | 0.4930 | 0.3907 | 0.4360 | 0.4070 |
| GLFM-VPDT-Pointwise | 0.0073 | 0.0036 | 0.0124 | 0.0124 | 0.0041 | 0.0046 | 0.0064 |
| LFM-V-Logit | 0.4842 | 0.4105 | 0.2585 | 0.4163 | 0.2830 | 0.3126 | 0.4565 |
| LFM-T-Logit | 0.2684 | 0.2621 | 0.2274 | 0.2965 | 0.2375 | 0.2415 | 0.2599 |
| LFM-P-Logit | 0.1977 | 0.1805 | 0.2030 | 0.2161 | 0.2018 | 0.2086 | 0.1931 |
| LFM-D-Logit | 0.1505 | 0.1499 | 0.1906 | 0.2044 | 0.2110 | 0.2286 | 0.1470 |
| GLFM-Logit | 0.7073 | 0.5673 | 0.5200 | 0.6797 | 0.4290 | 0.4738 | 0.6801 |
| GLFM-V-Logit | 0.7431 | 0.6026 | 0.5242 | 0.6828 | 0.4550 | 0.5026 | 0.7176 |
| GLFM-VPD-Logit | 0.7200 | 0.5894 | 0.5081 | 0.6763 | 0.4388 | 0.4847 | 0.6963 |
| GLFM-VPDT-Logit | 0.6463 | 0.5742 | 0.3882 | 0.5870 | 0.3849 | 0.4251 | 0.6344 |
| LFM-V-Probit | 0.4313 | 0.3708 | 0.2287 | 0.3391 | 0.2531 | 0.2824 | 0.4011 |
| LFM-T-Probit | 0.2836 | 0.2773 | 0.2289 | 0.2315 | 0.2410 | 0.2454 | 0.2745 |
| LFM-P-Probit | 0.2018 | 0.1991 | 0.1742 | 0.1855 | 0.1752 | 0.1863 | 0.1984 |
| LFM-D-Probit | 0.1712 | 0.1616 | 0.1542 | 0.1747 | 0.2015 | 0.2044 | 0.1675 |
| GLFM-Probit | 0.7136 | 0.5678 | 0.5328 | 0.6748 | 0.4378 | 0.4835 | 0.6879 |
| GLFM-V-Probit | 0.7221 | 0.6047 | 0.4978 | 0.6801 | 0.4396 | 0.4855 | 0.7021 |
| GLFM-VPD-Probit | 0.7294 | 0.6042 | 0.4968 | 0.6685 | 0.4464 | 0.4930 | 0.7082 |
| GLFM-VPDT-Probit | 0.6842 | 0.5842 | 0.4368 | 0.6083 | 0.4146 | 0.4578 | 0.6683 |

- The Logistic and Probit pairwise probability functions yield very competitive results while the Logistic function generates slightly better performance than the Probit function does. All the best results on San Francisco and Washington DC are attained by the Logistic function. For New York and Chicago, the results are mixed.
- Among the baselines, the Group-Membership method obtains the best performance on all the four cities, which further validates our assumption that group is an important factor for event recommendation. On the other hand, there still

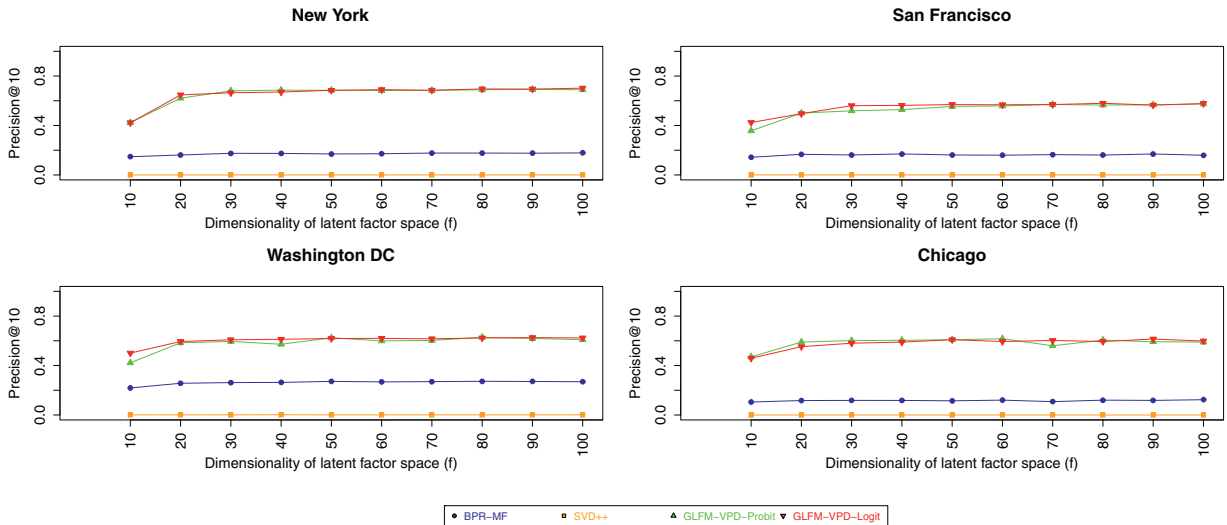


Fig. 2. Effect of dimensionality of latent factor space for P@10.

exists a large gap between the results of Group-Membership and the proposed group-aware models, which demonstrates the effectiveness of our dual perspective of group information.

- Both User-KNN and Item-KNN performed better than Item Mean, Biased-MF and SVD++ models. These results are consistent with Macedo et al. (2015) which found that state-of-the-art matrix factorization algorithms did not perform better than neighborhood-based methods in event recommendation. In fact, the pointwise variation of the our model, GLFM-VPD-Pointwise, also yields poor results that are similar to Biased-MF and SVD++. The BPR-MF model, on the other hand, yields the best results among the matrix factorization based baselines, validating our decision to utilize the pairwise ranking approach for event recommendation. SVDFeature also generates good results, which are second only to the BPR-MF model.

In sum, the experimental results consistently demonstrate the effectiveness of the proposed latent factor models by exploiting the dual-perceptive of group information with pairwise learning.

5.3.2. Effect of dimensionality of latent factor space

In this section, we investigate the effect of the number of latent factors f (i.e., dimensionality of the latent factor space) on our proposed models (GLFM-VPD-Logit and GLFM-VPD-Probit) and other state-of-the-art latent factor models (BPR-MF and SVD++) in event recommendation. The number of latent factors is varied from 10 to 100 in increments of 10. Fig. 2 plots the results in P@10 for the four cities, respectively. We have the following observations from the figure:

- The proposed group-aware latent factor models (GLFM-VPD-Logit and GLFM-VPD-Probit) perform significantly better than the other latent factor methods at most values of f for all the cities. These results demonstrate consistent improvement of the proposed models over the baselines across different numbers of latent factors.
- The results of the group-aware latent factor models gradually improve until $f = 40$ and plateau after that with no significant improvement. This pattern is observed for all four cities. On the other hand, the results of GLFM-VPD-Logit and GLFM-VPD-Probit are quite similar at different f while GLFM-VPD-Logit yields slightly better performance than GLFM-VPD-Probit in the majority of the cases of f .
- The results of the BPR-MF method are significantly better than SVD++ across all values of f while there is not much variation in the results of the two methods as the value of f increases from 10 to 100. The same pattern is observed for all four cities. These results validate the advantage of using the pairwise ranking approach. A noticeable observation is that the recall results of BPR-MF for all four cities are almost similar to the results for the proposed group-aware latent factor model (GLFM-VPD-Logit and GLFM-VPD-Probit) for $f = 10$ to 20, but the group-aware latent factor models perform significantly better on recall for values of $f > 30$.

As we can see, the relative performances of all the latent-factor based methods seem quite stable. These results indicate the dimensionality of the latent factor space may not be very sensitive. We set $f = 40$ as the default value in other experiments.

5.3.3. Preference instance generation

Section 4.6 introduces how preference instances are derived based on three relations. In this section, we compare two strategies for generating the preference instances. The first strategy is the default one shown in Section 4.6, which we

Table 10

Comparison between the default preference instance generation strategy WUP and the alternative strategy WOUP with the GLFM-VPD-Logit model for Chicago.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|-----------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| GLFM-VPD-Logit (WOUP) | 0.1989 | 0.1947 | 0.1260 | 0.2086 | 0.1093 | 0.1208 | 0.1965 |
| GLFM-VPD-Logit (WUP) | 0.7200 | 0.5894 | 0.5081 | 0.6763 | 0.4388 | 0.4847 | 0.6963 |

Table 11

Experimental results in the cold-start setting on New York.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| New Users | | | | | | | |
| Item Mean | 0.0031 | 0.0021 | 0.0139 | 0.0194 | 0.0088 | 0.0096 | 0.0029 |
| Group-Membership | 0.0003 | 0.0003 | 0.0018 | 0.0036 | 0.0033 | 0.0036 | 0.0003 |
| Biased-MF | 0.0001 | 0.0001 | 0.0005 | 0.0011 | 0.0004 | 0.0005 | 0.0003 |
| BPR-MF | 0.0163 | 0.0124 | 0.0769 | 0.1178 | 0.0219 | 0.0266 | 0.0202 |
| SVD++ | 0.0001 | 0.0001 | 0.0005 | 0.0010 | 0.0005 | 0.0005 | 0.0003 |
| SVDFeature | 0.0124 | 0.0118 | 0.0547 | 0.0919 | 0.0183 | 0.0191 | 0.0116 |
| GLFM-VPD-Logit | 0.0477 | 0.0293 | 0.2107 | 0.2618 | 0.0424 | 0.0468 | 0.0579 |
| GLFM-VPD-Probit | 0.0486 | 0.0280 | 0.2163 | 0.2507 | 0.0407 | 0.0450 | 0.0562 |
| New Events | | | | | | | |
| Group-Membership | 0.0025 | 0.0025 | 0.0027 | 0.0053 | 0.0039 | 0.0044 | 0.0025 |
| Biased-MF | 0.0001 | 0.0001 | 0.0002 | 0.0009 | 0.0003 | 0.0003 | 0.0002 |
| BPR-MF | 0.0156 | 0.0121 | 0.0721 | 0.1118 | 0.0317 | 0.0331 | 0.0275 |
| SVD++ | 0.0001 | 0.0001 | 0.0005 | 0.0011 | 0.0005 | 0.0005 | 0.0003 |
| SVDFeature | 0.0108 | 0.0100 | 0.0523 | 0.0667 | 0.0204 | 0.0202 | 0.0102 |
| GLFM-VPD-Logit | 0.0509 | 0.0292 | 0.2263 | 0.2619 | 0.0432 | 0.0477 | 0.0592 |
| GLFM-VPD-Probit | 0.0463 | 0.0291 | 0.2053 | 0.2619 | 0.0394 | 0.0435 | 0.0551 |

denote as WUP (With Unknown RSVPs). It forms a preference pair with an unknown RSVP as the positive instance and a random negative RSVP as the negative instance. The second strategy, which we denotes as WOUP, does not assume that the unknown RSVPs are preferred over the observed RSVP “No”. The results of both the strategies are provided in Table 10 (only the GLFM-VPD-Logit model for Chicago is shown to avoid clutter since GLFM-VPD-Logit usually gives the best results across different cities as shown in the previous sections). As we can see, the default strategy WUP yields much superior results than WOUP does, which validates the advantage of utilizing the unobserved RSVPs.

5.3.4. Cold-start event recommendation

Cold-start is a challenging problem in any recommendation system when there is no information about users or items. Cold-start is especially prevalent in event recommendation, because events are always in the future and short-lived. In this section, we test the performance of our models in the cold-start settings. We evaluate our models under two scenarios: *new users* and *new items*.

For the *new users* scenario, we split the dataset to ensure the same user is not present in both the training and test sets. The time order is reserved so that the events in the test data always occur after those in the training data. Similar to what we did for the experiments in the regular setting, we perform a training-test data split of 80% vs. 20%. We split the dataset similarly for the *new events* scenario with a 80% vs. 20% ratio for training and testing sets and ensure that the same event is not present in both sets.

We perform the experiments on the group-aware models that consider the group, venue, popularity and distance factors (*GLFM-VPD-Logit* and *GLFM-VPD-Probit*) as they provide the best results for a majority of the regular experiments. We compare the performance of group-aware latent factor models with the baseline methods: Item Mean, Group-Membership, Biased-MF, BPR-MF, SVD++, and SVDFeature. Item Mean is not applicable to the *new events* scenario due to no historical information available for the new events. UserKNN and ItemKNN cannot handle the cold-start settings and hence we ignore them in this experiment. Regarding our proposed models, since a user in the test set is not present in the training set, we are unable to learn the user-specific parameters. In other words, we do not have user latent factor \mathbf{p}_u , but we have the dual-perspective group factors \mathbf{r}_g and \mathbf{t}_g that enable the calculation of the ranking score for the *new users* scenario. The prediction of our models is made by removing p_u from Eq. (11). Similarly, for the *new items* scenario, we do not have event latent factor \mathbf{q}_i and the popularity count c_i , but we have the event-oriented factor \mathbf{t}_g . The prediction is made by removing \mathbf{q}_i and c_i from Eq. (11).

Tables 11–14 contain the results for the cold-start experiments. As we can see, the best results in both new-user and new-event scenarios are achieved by our proposed methods, *GLFM-VPD-Logit* or *GLFM-VPD-Probit*. This pattern holds true across different cities and different evaluation metrics. The improvement is especially substantial on New York and Washington DC. Furthermore, *GLFM-VPD-Logit* and *GLFM-VPD-Probit* yield competitive results. *GLFM-VPD-Probit* achieves the best results on

Table 12
Experimental results in the cold-start setting on San Francisco.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| New Users | | | | | | | |
| Item Mean | 0.0072 | 0.0074 | 0.0320 | 0.0688 | 0.0091 | 0.0108 | 0.0084 |
| Group-Membership | 0.0003 | 0.0003 | 0.0016 | 0.0033 | 0.0029 | 0.0030 | 0.0003 |
| Biased-MF | 0.0001 | 0.0001 | 0.0002 | 0.0005 | 0.0002 | 0.0002 | 0.0001 |
| BPR-MF | 0.0270 | 0.0180 | 0.1291 | 0.1526 | 0.0177 | 0.0189 | 0.0238 |
| SVD++ | 0.0001 | 0.0001 | 0.0003 | 0.0011 | 0.0003 | 0.0004 | 0.0002 |
| SVDFeature | 0.0194 | 0.0187 | 0.1011 | 0.1276 | 0.0163 | 0.0171 | 0.0190 |
| GLFM-VPD-Logit | 0.0201 | 0.0122 | 0.0765 | 0.0947 | 0.0180 | 0.0199 | 0.0235 |
| GLFM-VPD-Probit | 0.0308 | 0.0186 | 0.1299 | 0.1590 | 0.0242 | 0.0267 | 0.0335 |
| New Events | | | | | | | |
| Group-Membership | 0.0075 | 0.0071 | 0.0031 | 0.0059 | 0.0042 | 0.0049 | 0.0075 |
| Biased-MF | 0.0001 | 0.0001 | 0.0003 | 0.0009 | 0.0002 | 0.0003 | 0.0001 |
| BPR-MF | 0.0270 | 0.0181 | 0.1096 | 0.1171 | 0.0093 | 0.0105 | 0.0826 |
| SVD++ | 0.0001 | 0.0001 | 0.0003 | 0.0010 | 0.0003 | 0.0004 | 0.0002 |
| SVDFeature | 0.0186 | 0.0173 | 0.1144 | 0.1261 | 0.0179 | 0.0185 | 0.0164 |
| GLFM-VPD-Logit | 0.0171 | 0.0114 | 0.0647 | 0.0887 | 0.0145 | 0.0160 | 0.0191 |
| GLFM-VPD-Probit | 0.0281 | 0.0185 | 0.1183 | 0.1549 | 0.0200 | 0.0221 | 0.0297 |

Table 13
Experimental results in the cold-start setting on Washington DC.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| New Users | | | | | | | |
| Item Mean | 0.0013 | 0.0012 | 0.0068 | 0.0101 | 0.0084 | 0.0097 | 0.0011 |
| Group-Membership | 0.0004 | 0.0004 | 0.0019 | 0.0037 | 0.0024 | 0.0026 | 0.0004 |
| Biased-MF | 0.0001 | 0.0001 | 0.0006 | 0.0012 | 0.0005 | 0.0006 | 0.0004 |
| BPR-MF | 0.0150 | 0.0094 | 0.0143 | 0.0186 | 0.0107 | 0.0119 | 0.0147 |
| SVD++ | 0.0001 | 0.0001 | 0.0008 | 0.0020 | 0.0007 | 0.0009 | 0.0005 |
| SVDFeature | 0.0124 | 0.0119 | 0.0167 | 0.0186 | 0.0144 | 0.0151 | 0.0111 |
| GLFM-VPD-Logit | 0.0165 | 0.0108 | 0.0671 | 0.0870 | 0.0127 | 0.0129 | 0.0184 |
| GLFM-VPD-Probit | 0.0174 | 0.0126 | 0.0733 | 0.1054 | 0.0133 | 0.0147 | 0.0197 |
| New Events | | | | | | | |
| Group-Membership | 0.0004 | 0.0004 | 0.0015 | 0.0032 | 0.0026 | 0.0029 | 0.0004 |
| Biased-MF | 0.0002 | 0.0001 | 0.0009 | 0.0015 | 0.0006 | 0.0008 | 0.0005 |
| BPR-MF | 0.0144 | 0.0101 | 0.0207 | 0.0261 | 0.0094 | 0.0107 | 0.0188 |
| SVD++ | 0.0002 | 0.0001 | 0.0009 | 0.0017 | 0.0009 | 0.0009 | 0.0006 |
| SVDFeature | 0.0131 | 0.0118 | 0.0219 | 0.0221 | 0.0171 | 0.0188 | 0.0126 |
| GLFM-VPD-Logit | 0.0180 | 0.0109 | 0.0784 | 0.0935 | 0.0122 | 0.0134 | 0.0194 |
| GLFM-VPD-Probit | 0.0173 | 0.0111 | 0.0688 | 0.0903 | 0.0096 | 0.0106 | 0.0161 |

Table 14
Experimental results in the cold-start setting on Chicago.

| Method | P@5 | P@10 | R@5 | R@10 | NDCG@5 | NDCG@10 | MAP@10 |
|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| New Users | | | | | | | |
| Item Mean | 0.0011 | 0.0019 | 0.0045 | 0.0164 | 0.0041 | 0.0057 | 0.0014 |
| Group-Membership | 0.0005 | 0.0005 | 0.0019 | 0.0032 | 0.0037 | 0.0039 | 0.0005 |
| Biased-MF | 0.0002 | 0.0002 | 0.0010 | 0.0023 | 0.0008 | 0.0010 | 0.0006 |
| BPR-MF | 0.0234 | 0.0181 | 0.1073 | 0.1662 | 0.0120 | 0.0153 | 0.0296 |
| SVD++ | 0.0003 | 0.0003 | 0.0015 | 0.0036 | 0.0016 | 0.0017 | 0.0008 |
| SVDFeature | 0.0194 | 0.0182 | 0.0997 | 0.1055 | 0.0193 | 0.0217 | 0.0187 |
| GLFM-VPD-Logit | 0.0276 | 0.0203 | 0.1105 | 0.1730 | 0.0204 | 0.0225 | 0.0318 |
| GLFM-VPD-Probit | 0.0190 | 0.0146 | 0.0779 | 0.1200 | 0.0158 | 0.0175 | 0.0234 |
| New Events | | | | | | | |
| Group-Membership | 0.0005 | 0.0005 | 0.0017 | 0.0034 | 0.0042 | 0.0061 | 0.0005 |
| Biased-MF | 0.0003 | 0.0003 | 0.0015 | 0.0034 | 0.0012 | 0.0015 | 0.0009 |
| BPR-MF | 0.0176 | 0.0158 | 0.0822 | 0.1492 | 0.0132 | 0.0146 | 0.0199 |
| SVD++ | 0.0005 | 0.0003 | 0.0022 | 0.0034 | 0.0011 | 0.0016 | 0.0007 |
| SVDFeature | 0.0167 | 0.0158 | 0.0822 | 0.0945 | 0.0171 | 0.0183 | 0.0159 |
| GLFM-VPD-Logit | 0.0311 | 0.0199 | 0.1308 | 0.1703 | 0.0199 | 0.0220 | 0.0312 |
| GLFM-VPD-Probit | 0.0186 | 0.0173 | 0.07403 | 0.1470 | 0.0132 | 0.0146 | 0.0220 |

San Francisco while *GLFM-VPD-Logit* obtains the best on Chicago for both scenarios. On the other two cities, the results are mixed. Among the baseline methods, BPR-MF generates the best results across four cities in all the metrics for both cold-start scenarios. These results validate the advantage of pairwise training for the event recommendation task. In sum, the experimental results demonstrate the advantage of the proposed models in dealing with the cold-start problems for event recommendation. In our cold-start experiments, we considered *only* new users and events, which is atypical of a real-world scenario.

6. Conclusions and future work

We systematically investigate the effect of group information on event recommendation. A latent factor model is proposed based on the dual-perspective of groups. Logistic and Probit functions are used to model the probability of pairwise preferences that consist of observed and unobserved user feedback. Additional contextual information such as event venue, popularity, temporal influence, and geographical distance can be readily incorporated into the model. The experiments on the Meetup data of four cities demonstrate the importance of group information and show much improved performance over the state-of-the-art baselines. Moreover, the proposed approach demonstrates advantages of tackling the cold-start problems by utilizing the dual role of groups.

The proposed dual-perspective latent factor model can be applied to other recommendation tasks where certain factors may have a dual view. For example, users may specify their topics of interest in their profiles (e.g., music recommendation in Pandora, job recommendation in LinkedIn, book recommendation in Amazon, etc.) and items may also have the topic information available (e.g., genre of a song and category of a job or book). In this case, the topics serve a dual role. In future work, we will generalize the proposed approach to a wide range of recommendation tasks where a dual perspective of factors is present. In addition, we plan to incorporate into the proposed model more contextual features such as content information (Du et al., 2014; Macedo et al., 2015; Zhang & Wang, 2015), and social relations (Boutsis, Karanikolaou, & Kalogeraki, 2015; Qiao et al., 2014a), to further boost the performance of event recommendation. Moreover, we will investigate learning to rank based recommendation (Belem, Martins, Almeida, & Goncalves, 2014) such as the listwise recommendation approach by taking a ranked list of items as a training instance to minimize a loss function over the training ranked lists. A listwise recommendation approach has demonstrated effectiveness and efficiency over the pairwise methods in movie recommendations (Shi, Larson, & Hanjalic, 2010). Last but not the least, we will take a deep look into the latent factors learned by the proposed models and explore methods to make the results more explainable.

References

- Agarwal, D., & Chen, B.-C. (2009). Regression-based latent factor models. In *Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 19–28). ACM.
- Agresti, A., & Kateri, M. (2014). Some remarks on latent variable models in categorical data analysis. *Communications in Statistics-Theory and Methods*, 43(4), 801–814.
- Belem, F. M., Martins, E. F., Almeida, J. M., & Goncalves, M. A. (2014). Personalized and object-centered tag recommendation methods for web 2.0 applications. *Information Processing & Management*, 50(4), 524–553.
- Bell, R., Koren, Y., & Volinsky, C. (2007). Modeling relationships at multiple scales to improve accuracy of large recommender systems. In *Proceedings of the 13th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 95–104). ACM.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *The Journal of Machine Learning Research*, 3, 993–1022.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT* (pp. 177–186). Springer.
- Boutsis, I., Karanikolaou, S., & Kalogeraki, V. (2015). Personalized event recommendations using social networks. In *Proceedings of the IEEE international conference on mobile data management: Vol. 1* (pp. 84–93). IEEE.
- Breese, J. S., Heckerman, D., & Kadie, C. (1998). Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the fourteenth conference on uncertainty in artificial intelligence* (pp. 43–52). Morgan Kaufmann Publishers Inc..
- Cao, L. (2015). Coupling learning of complex interactions. *Information Processing & Management*, 51(2), 167–186.
- Chen, C. C., & Sun, Y.-C. (2016). Exploring acquaintances of social network site users for effective social event recommendations. *Information Processing Letters*, 116(3), 227–236.
- Chen, T., Zhang, W., Lu, Q., Chen, K., Zheng, Z., & Yu, Y. (2012). Svdfeature: A toolkit for feature-based collaborative filtering. *The Journal of Machine Learning Research*, 13(1), 3619–3622.
- Du, R., Yu, Z., Mei, T., Wang, Z., Wang, Z., & Guo, B. (2014). Predicting activity attendance in event-based social networks: Content, context and social influence. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing* (pp. 425–434). ACM.
- Formoso, V., FernáNdez, D., Cacheda, F., & Carneiro, V. (2013). Using profile expansion techniques to alleviate the new user problem. *Information Processing & Management*, 49(3), 659–672.
- Gantner, Z., Drumond, L., Freudenthaler, C., & Schmidt-Thieme, L. (2012). Personalized ranking for non-uniformly sampled items. In *KDD cup* (pp. 231–247).
- Hu, L., Sun, A., & Liu, Y. (2014). Your neighbors affect your ratings: on geographical neighborhood influence to rating prediction. In *Proceedings of the 37th international ACM SIGIR conference on research & development in information retrieval* (pp. 345–354). ACM.
- Jiang, J.-Y., & Li, C.-T. (2016). Analyzing social event participants for a single organizer. In *Tenth international AAAI conference on web and social media* (pp. 599–602).
- Kassak, O., Kompan, M., & Bielikova, M. (2016). Personalized hybrid recommendation for group of users: Top-n multimedia recommender. *Information Processing & Management*, 52(3), 459–477.
- Khrouf, H., & Troncy, R. (2013). Hybrid event recommendation using linked data and user diversity. In *Proceedings of the 7th ACM conference on recommender systems* (pp. 185–192). ACM.
- Koren, Y. (2008). Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 426–434). ACM.
- Koren, Y. (2010). Collaborative filtering with temporal dynamics. *Communications of the ACM*, 53(4), 89–97.
- Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37.
- Krohn-Grimberghe, A., Drumond, L., Freudenthaler, C., & Schmidt-Thieme, L. (2012). Multi-relational matrix factorization using Bayesian personalized ranking for social network data. In *Proceedings of the fifth ACM international conference on web search and data mining* (pp. 173–182). ACM.

- Li, X., Cong, G., Li, X.-L., Pham, T.-A. N., & Krishnaswamy, S. (2015). Rank-geofm: A ranking based geographical factorization method for point of interest recommendation. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval* (pp. 433–442). ACM.
- Liao, S.-h., & Chang, H.-k. (2016). A rough set-based association rule approach for a recommendation system for online consumers. *Information Processing & Management, in press*.
- Liu, N. N., & Yang, Q. (2008). Eigenrank: A ranking-oriented approach to collaborative filtering. In *Proceedings of the 31st annual international ACM SIGIR conference on research and development in information retrieval* (pp. 83–90). ACM.
- Liu, N. N., Zhao, M., & Yang, Q. (2009). Probabilistic latent preference analysis for collaborative filtering. In *Proceedings of the 18th ACM conference on information and knowledge management* (pp. 759–766). ACM.
- Liu, X., He, Q., Tian, Y., Lee, W.-C., McPherson, J., & Han, J. (2012). Event-based social networks: linking the online and offline social worlds. In *Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1032–1040). ACM.
- Lu, D., Voss, C., Tao, F., Ren, X., Guan, R., Korolov, R., ... Kaplan, L. (2016). Cross-media event extraction and recommendation. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Demonstrations* (pp. 72–76). San Diego, California: Association for Computational Linguistics.
- de Macedo, A. Q., & Marinho, L. B. (2014). Event recommendation in event-based social networks. In *Proceedings of the 25th ACM hypertext and social media conference - social personalization workshop* (pp. 7–12).
- Macedo, A. Q., Marinho, L. B., & Santos, R. L. (2015). Context-aware event recommendation in event-based social networks. In *Proceedings of the 9th ACM conference on recommender systems* (pp. 123–130). ACM.
- Manning, C. D., Raghavan, P., Schütze, H., et al. (2008). *Introduction to information retrieval: Vol. 1*. Cambridge university press Cambridge.
- Martins, E. F., Belém, F. M., Almeida, J. M., & Gonçalves, M. A. (2016). On cold start for associative tag recommendation. *Journal of the Association for Information Science and Technology*, 67(1), 83–105.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models: 37*. CRC press.
- Minkov, E., Charrow, B., Ledlie, J., Teller, S., & Jaakkola, T. (2010). Collaborative future event recommendation. In *Proceedings of the 19th ACM international conference on information and knowledge management* (pp. 819–828). ACM.
- Mnih, A., & Salakhutdinov, R. (2007). Probabilistic matrix factorization. In *Advances in neural information processing systems* (pp. 1257–1264).
- Pan, W., & Chen, L. (2013). Gbpr: Group preference based Bayesian personalized ranking for one-class collaborative filtering. In *IJCAI: 13* (pp. 2691–2697).
- Pazzani, M. J., & Billsus, D. (2007). Content-based recommendation systems. In *The adaptive web* (pp. 325–341). Springer.
- Pera, M. S., & Ng, Y.-K. (2013). A group recommender for movies based on content similarity and popularity. *Information Processing & Management*, 49(3), 673–687.
- Pham, T.-A. N., Li, X., Cong, G., & Zhang, Z. (2015). A general graph-based model for recommendation in event-based social networks. In *Proceedings of IEEE international conference on data engineering* (pp. 567–578). IEEE.
- Qiao, Z., Zhang, P., Cao, Y., Zhou, C., Guo, L., & Fang, B. (2014a). Combining heterogenous social and geographical information for event recommendation. In *Proceedings of the twenty-eighth AAAI conference on artificial intelligence* (pp. 145–151).
- Qiao, Z., Zhang, P., Zhou, C., Cao, Y., Guo, L., & Zhang, Y. (2014b). Event recommendation in event-based social networks. In *Proceedings of the twenty-eighth AAAI conference on artificial intelligence* (pp. 3130–3131).
- Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2009). Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence* (pp. 452–461). AUAI Press.
- Rendle, S., & Schmidt-Thieme, L. (2010). Pairwise interaction tensor factorization for personalized tag recommendation. In *Proceedings of the third ACM international conference on web search and data mining* (pp. 81–90). ACM.
- 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* (pp. 285–295). ACM.
- Shi, Y., Larson, M., & Hanjalic, A. (2010). List-wise learning to rank with matrix factorization for collaborative filtering. In *Proceedings of the fourth ACM conference on recommender systems* (pp. 269–272). ACM.
- Shi, Y., Zhao, X., Wang, J., Larson, M., & Hanjalic, A. (2012). Adaptive diversification of recommendation results via latent factor portfolio. In *Proceedings of the 35th international ACM SIGIR conference on research and development in information retrieval* (pp. 175–184). ACM.
- Shumaker, B., & Sinnott, R. (1984). *Astronomical computing: Vol. 1. Computing under the open sky. 2. Virtues of the haversine.* *Sky and Telescope*, 68, 158–159.
- Singh, A. P., & Gordon, G. J. (2008). Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 650–658). ACM.
- Sun, J., Wang, G., Cheng, X., & Fu, Y. (2015). Mining affective text to improve social media item recommendation. *Information Processing & Management*, 51(4), 444–457.
- Wang, C., & Blei, D. M. (2011). Collaborative topic modeling for recommending scientific articles. In *Proceedings of the 17th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 448–456). ACM.
- Yao, L., Sheng, Q. Z., Qin, Y., Wang, X., Shemshadi, A., & He, Q. (2015). Context-aware point-of-interest recommendation using tensor factorization with social regularization. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval* (pp. 1007–1010). ACM.
- Zhang, W., & Wang, J. (2015). A collective Bayesian poisson factorization model for cold-start local event recommendation. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1455–1464). ACM.
- Zhang, W., Wang, J., & Feng, W. (2013). Combining latent factor model with location features for event-based group recommendation. In *Proceedings of the 19th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 910–918). ACM.
- Zhao, T., McAuley, J., & King, I. (2015). Improving latent factor models via personalized feature projection for one class recommendation. In *Proceedings of the 24th ACM international on conference on information and knowledge management* (pp. 821–830). ACM.