HBGG: a Hierarchical Bayesian Geographical Model for Group Recommendation

Ziyu Lu* Hui Li* Nikos Mamoulis* David W. Cheung*

Abstract

Location-based social networks such as Foursquare and Plancast have gained increasing popularity. On those sites, users can organize and participate in group activities; hence, recommending venues to a group is of practical importance. In this paper, we study the problem of recommending venues to groups of users and propose a Hierarchical Bayesian Model (HBGG) for this purpose. First, a generative group geographical topic model (GG) which exploits group membership, group mobility regions and group preferences is proposed. And we integrate social structure into one-class collaborative filtering as social-based collaborative filtering (SOCF) to leverage social wisdom. Through the shared latent group features, HBGG connects the group geographical model with SOCF framework for group recommendation. Experimental results on two real datasets show that our methods outperforms the state-of-the-art group recommenders, especially on cold-start user groups.

1 Introduction

With the development of location-based social networking services such as Facebook[1], Meetup[2] and Plancast[3], people can easily organize and participate in group activities. In addition, the participants of such services typically enjoy sharing their social connections and group activities online. Given the above, recommending products or activities to a group of users has become an important service. On the other hand, group recommendation is much more challenging compared to personalized recommendation offered by traditional recommender systems [1, 24, 14]. The main difficulty in group recommendation is to understand how group decisions are made, especially for an ad-hoc group consisting of multiple users with different preferences.

Early studies on group recommendation [3, 2] aggregate recommendations to individual group members. These approaches are shown to be inadequate because they ignore the influences between users in the group decision. Recently, several model-based group recommendation methods have been proposed [23, 11, 26], which consider social influences. However, although location information has been shown to play an important role in location recommendation [12, 25], past studies on group recommendation do not explicitly consider the problem of venue recommendation to a group and do not model the geographical influences of users. Specifically, there are no studies to-date which consider *group mobility behaviors* in group recommendation. Besides, [16] mentioned that the data sparsity and cold-start problems are very severe in group recommendation. Collaborative filtering has been used to alleviate data sparsity. Group check-ins are implicit data that can help, however, they do not carry negative feedback, i.e., check-ins do not indicate whether a user dislikes a location, and check-in frequency is not a reliable indicator of how much a user likes a location. Recently, generating recommendations for implicit data has attracted increasing attention, leading to techniques such as one-class collaborative filtering (OCCF) [6, 15] and Bayesian Personalized Ranking (BPR) [18].

In this paper, we first design a generative group geographical (GG) topic model to capture group preferences from group membership and group geographical regions. Then, we integrate social structure into one-class collaborative filtering as social-based collaborative filtering (SOCF) to explore social influences on making group decisions for group recommendation. We combine GG and SOCF models into a Hierarchical Bayesian model via the shared group latent features. The group geographical topic model interprets group latent features from group member engagement and group mobility regions. The social-based collaborative filtering part can further enhance the learning of group latent features from social wisdom, which can help to address issues of cold-start groups. The contributions of our work are listed as follows:

- We utilize group mobility regions for group events, and propose a generative group geographical topic model (GG) for group recommendation.
- We integrate social connections into one-class collaborative filtering as social-based collaborative filtering (SOCF) to address the cold-start problem.
- We propose a *Hierarchical Bayesian Geographical* model (HBGG) which combines GG with SOCF through a shared group latent features.
- We present experiments on two real datasets (Foursquare and Plancast), showing that our proposed methods outperform the state-of-the-art methods by a wide margin.

The rest of the paper is organized as follows. Section 2

*Department of Computer Science, The University of Hong Kong, Hong Kong. {zylu,hli2,nikos,dcheung}@cs.hku.hk

[1]https://www.facebook.com/places

[2]http://www.meetup.com/

[3]http://plancast.com

presents related work, followed by problem definition and background in Section 3. Section 4 presents the proposed models. Section 5 reports the experimental results. Finally we conclude our paper in Section 6.

2 Related Work

2.1 Recommender Systems Recommender systems for individual users have been extensively studied in the past two decades. There are three types of recommendation methods: *content-based* approaches, *collaborative filtering* (CF) approaches and *hybrid* methods, which combine content-based recommendation and collaborative filtering [1]. Content-based approaches extract features from the items (e.g., keywords); candidate items are compared with items previously selected by the user and the most similar ones in the feature space are recommended. Collaborative filtering leverages the history records of users in order to generate recommendations. CF methods can be divided into two categories: memory-based and model-based techniques. Memory-based methods include user-based CF and item-based CF [20]. User-based CF finds the most similar users to the target user, by comparing their vectors of item ratings using some similarity measure (e.g., cosine similarity). The predicted score of an unrated item by the target user is the similarity-weighted average of other users' preferences on the item. Item-based CF [20] takes a transposed view of user-based CF. It firstly computes similarity between items and the recommendation score of an item is the similarity-weighted average of ratings given by the target user on similar items. Model-based CF make recommendations based on learning models such as probabilistic topic models [24, 12] and latent factor models that employ matrix factorization [7, 19]. Recently, CF methods for implicit data have been proposed such as one-class collaborative filtering (OCCF) [6, 15] and Bayesian Personalized Ranking (BPR) [18].

2.2 Group Recommendation Early studies on group recommendation extend memory-based CF, based on aggregation methods [3, 2], including *preference aggregation* and *score aggregation*. The difference between the two is the order of applying the aggregation and recommendation steps. *Preference aggregation* methods first aggregate the profiles of all group members and then perform recommendation on the aggregated profile using a classic CF approach (e.g., user-based CF). *Score-based aggregation* approaches firstly compute recommendations for each group member and then aggregate them. Aggregation is done by taking for each item the *average* (averaging) or *minimum* (least-misery) predicted score per user and considering this to be the item's score for the group. In addition, the diversity of members' preferences can be integrated with the aggregated preferences to construct a consensus function for group recommendation [2]. The main drawback of aggregation-based methods is that they ignore the influences between users in the group decision.

Several model-based methods have been proposed for group recommendation. Social influences [23, 11, 8, 9] have been explored. For example, Liu et al. [11] proposed a personal impact topic (PIT) model for group recommendation and an extended model with social information. PIT builds an author-topic model and differentiates personal impacts for group decision by introducing a personal impact variable to represent the relative probability to influence the group's decision in the PIT model. PIT uses learned personal impacts to aggregate individual preference scores as group preference scores. In addition, Yuan et al. [26] proposed a generative model COM for group recommendation based on two different assumptions: (i) personal impacts are topic-dependent and (ii) the group decision process depends on the group's topic preferences and the personal preferences of individual group members. For the first assumption, COM uses the group-topic and topic-user distributions to represent the topic-dependent personal influence within the group. For the second assumption, COM uses a personal variable as a weight when combining the group's topic preferences and personal preferences. COM uses geographical and content information as priors in the model, but it does not directly model geographical influences. Geographical information is very important in recommendations for individual and has been extensively investigated in location recommendation problems [13, 25]. However, there are no studies to-date which consider *group mobility behaviors* in group recommendation. We investigate group mobility behaviors in a generative group geographical topic model for group recommendation.

Purushotham et al. [16] proposed a Collaborative Group Activity Recommendation (CGAR) system to jointly learn group and activity latent spaces. CGAR fuses topic modeling with matrix factorization to obtain a consistent and compact feature representation based on Collaborative Topic Regression (CTR) [22, 17]. This is the most similar work to ours. Although CGAR model learns group latent features and activity latent features from the semantics and collaborative wisdom, it does not consider group mobility regions and influences from social structure. Our proposed model can mimic the group-item selection process by investigating group membership, group mobility region in a generative geographical topic model and exploiting influences from the social structure in the one-class collaborative filtering framework to effectively handle issues of cold-start groups.

3 Preliminaries

3.1 Problem Definition We assume the availability of historical information about visits of user groups to venues. Such information can easily be tracked by group event planning sites such as Plancast, or inferred by records of user visits in combination with their social relationships. For example, if we know that two users are friends and they visit a given venue within a short time difference (e.g. one hour [16]), we can infer that they visit the venue as a group. Formally, we assume that we have a set of users U, a set of venues V, a set G of historical user groups (i.e., for each $g \in G$, $g \subseteq U$), and a collection of historical group user visits. Each venue $v \in V$ has its corresponding geographical information $l_v = \langle latitude, longtitude \rangle$. The selections of venues by groups is represented as a implicit matrix $X \in \mathbb{R}^{|G| \times |V|}$, whereby X_{ij} indicates that group g_i selects venue v_j. In addition, we assume the availability of a *social influence* matrix $S \in \mathbb{R}^{|G| \times |U|}$, which captures the social influences during the group decision process; S_{ij} models the social weight of a user u_j to group g_i. The social weight of u_j to g_i is percentage of users in g_i who are friends with u_j (i.e., have an explicit connection with u_j in the social graph), namely $S_{ij} = \frac{|F(u_j, g_i)|}{|g_i|}$, where $F(u_j, g_i)$ is the set of u_j's friends in g_i. We can now provide a formal definition for the problem as follows:

DEFINITION 1. *Given the set of groups G, the set of users U, the set of venues V and the implicit group-venue matrix X and a social implicit matrix S, recommend to an ad-hoc user group $g \subseteq U$ the top-N venues to visit.*

3.2 Background In matrix factorization, given a user-item matrix H, user and item features are represented in a shared latent space of low dimensionality K. Then, a user u_i is represented as a latent vector $p_i \in \mathbb{R}^K$ and an item v_j as a latent vector $q_j \in \mathbb{R}^K$. The preferences of a user u_i to an item v_j can be quantified by the inner product $\hat{h}_{ij} = p_i q_j^T$ of u_i's and v_j's latent vectors. In one-class collaborative filtering for implicit data [6, 15], each observation h_{ij} is represented by a binary variable y_{ij} and a confidence value c_{ij}. y_{ij} indicates the implicit response of user u_i to item v_j; it is set to 1 if u_i has rated or checked-in v_j (i.e., positive observation), and 0 otherwise (i.e., negative observation). c_{ij} measures the confidence of y_{ij}; a higher c_{ij} indicates that y_{ij} is more trustable. Given the observed matrix X, we learn the known latent vectors q and p by minimizing the regularized squared error loss with respect to $P = (p_i)_{i=1}^{I}$ and $Q = (q_i)_{i=1}^{J}$ as follows:

$$min_{p^*,q^*} = \sum_{ij} c_{ij}(y_{ij} - p_i q_j^T) + \lambda_p \sum_i ||p_i||^2 + \lambda_q \sum_j ||q_j||^2,$$

where λ_p and λ_q are regularization parameters. Collaborative filtering can be generated as a probabilistic model (PMF) [19] . Then, for each user u_i, the user latent vector p_i is assumed to be drawn from $\mathcal{N}(0, \lambda_p^{-1} I_K)$ and for each item v_j, the latent vector q_j is drawn from $\mathcal{N}(0, \lambda_q^{-1} I_K)$. The implicit response y_{ij} is generated by: $y_{ij} \sim \mathcal{N}(p_i q_j^T, c_{ij}^{-1})$. In Section 4, we propose a Hierarchical Bayesian model in which the probabilistic matrix factorization model for implicit data is integrated with a probabilistic geographical topic model (based on LDA) for group recommendation.

Table 1: Notations used in GG and HBGG

Symbols	Description
N, M	number of venues, number of users
K, R	number of topics, number of regions
G, U, V	group latent factors, user latent factors, venue latent factors
X, S	implicit group-venue matrix, social matrix
A_g	the set of user members for the group g
z, r	latent topic variable, latent region variable
g, u, v	group variable, user variable, venue variable
θ_g	distribution of topics specific to a group g
ψ_z	distribution of users specific to topic z
π_z	distribution of regions specific to topic z
ϕ_r	distribution of venues specific to region r
ρ_z	distribution of venues specific to topic z
μ_r, Σ_r	mean venue and covariance of region r
$\alpha, \beta, \eta, \gamma, \omega$	Dirichlet priors for $\theta, \psi, \pi, \phi, \rho$
$\lambda_g, \lambda_u, \lambda_v$	regularized parameters for latent factors G, U, V

4 Our Models

In this section, we describe our proposed generative Group Geographical (GG) and Hierarchical Bayesian Geographical (HBGG) models for group recommendation of venues. HBGG combines GG and social-network based collaborative filtering (SOCF), which integrates social structure into one-class collaborative filtering [6, 15]. The main notations used in this section are summarized in Table 1. Figures 4.1 and 4.2 show the graphical representations of the proposed GG and HBGG models, respectively. We can see that HBGG adds the social-based collaborative filtering part on X and S to learn the latent group factor G, venue latent factor V, and user latent factor U (the upper right part in Figure 4.2). We first describe the general idea of GG and then present HBGG.

4.1 Group Geographical model (GG) The group geographical topic model (GG) is built upon three assumptions. (i) Group activities are topic dependent and they influence group membership. Different topics might have different group participants; e.g., the group is a set of friends or professionals if the activity is a social gathering or an academic workshop, respectively. (ii) Different groups have different mobility regions and these regions are topic dependent; e.g. a family would go to restaurants near home, while a group of office colleagues would go to pubs near their workplaces. Finally, (iii) group activity venues are influenced by both group topic interests and group mobility regions. In a nutshell, when a group g_i decides to organize or participate in an event, g_i first selects the topic of the event z (e.g., a social gathering or an academic activity), according to the group-topic distribution θ_{g_i}. The selected topic z influences the participant members A_{g_i} based on each user's personal topic preferences ψ_z and the event venue region r upon the topic-region distribution π_z. Finally, the event venue v_{g_i} is chosen, based on both the topic-venue distribution ρ_z and the geographical regions ϕ_r characterized by a Gaussian distribution. The detailed generative process of GG is as follows:

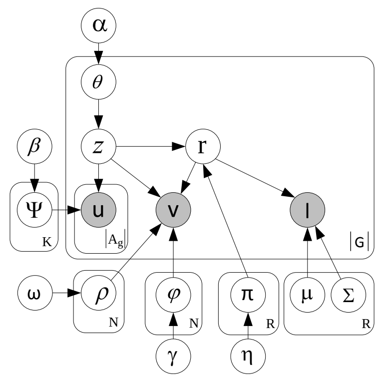

Figure 1: Graphical representation of GG

- For each topic z, draw ψ_z from $Dir(\beta)$.
- For each region r, draw π_r from $Dir(\eta)$.
- For each venue index v
  - Draw ϕ_v from $Dir(\gamma)$
  - Draw ρ_v from $Dir(\omega)$
- For each group g_i,
  - Draw θ_{g_i} from $Dir(\alpha)$
  - Draw a topic z from $Multi(\theta_{g_i})$
  - Draw a region r from $Multi(\pi_z)$
  - For each group member u in the member set A_{g_i}, draw u from $Multi(\psi_z)$
  - Draw a venue index v from a hybrid model $Multi(\phi_r) \times Multi(\rho_z)$
  - Draw a venue coordinate l_v from Gaussian Distribution $\mathcal{N}(\mu_r, \Sigma_r)$

4.2 Hierarchical Bayesian Geographical Model (HBGG) To further incorporate social influences into GG, we first apply one-class collaborative filtering (OCCF) using the social network of users to define a social-based collaborative filtering (SOCF) model. Then, SOCF and GG are combined together to form the Hierarchical Bayesian Geographical Model (HBGG).

SOCF interprets the implicit group item observations X and the group-user social connections S into group latent factors G, venue latent factors V and user latent factors U. The collaborative latent factors G interact with V and U. Then, given a new group (a cold-start group), the new group's preferences towards venues can be inferred by transferring preferences of existing groups using the social connections. The intuition behind this is that groups which have social connections might have common interests. In SOCF, X_{ij} is considered to be drawn from $\mathcal{N}(G_i V_j^T, w_{ij}^{-1})$ in which w_{ij} is the confidence parameter as introduced in Section 3.2. Similarly, the group-user social connections S_{im} are drawn from $\mathcal{N}(G_i U_m^T, d_{im}^{-1})$ and d_{im} is the confidence parameter for social connection.

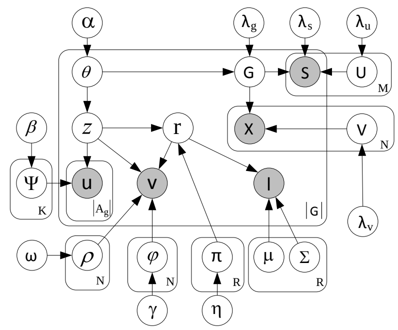

Figure 2: Graphical representation of HBGG

Then, HBGG combines GG (proposed in Section 4.1) with SOCF through the group latent factors G which are offset from the group-topic distribution θ. The introduced offset variable ϵ explains the degree to which the group decision making relies on the statistical inference from GG (offline social connections ψ and geographical influences π), and SOCF through online social relations. More specifically, the generative process of our Hierarchical Bayesian Geographical Model (HBGG) is as follows:

- For each user u, draw the user latent vector $U_m \sim \mathcal{N}(0, \lambda_u^{-1} I_K)$
- For each venue index v, draw the venue latent vector $V_j \sim \mathcal{N}(0, \lambda_v^{-1} I_K)$
- For each topic z, draw ψ_z from $Dir(\beta)$.
- For each region r, draw π_r from $Dir(\eta)$.
- For each venue index v
  - Draw ϕ_v from $Dir(\gamma)$,ρ_v from $Dir(\omega)$
- For each group g_i,
  - Draw θ_{g_i} from $Dir(\alpha)$
  - Draw group latent offset $\epsilon_{g_i} \sim \mathcal{N}(0, \lambda_{g_i}^{-1} I_K)$ and set the group latent vector as $G_i = \epsilon_{g_i} + \theta_{g_i}$
  - Draw a topic z from $Multi(\theta_{g_i})$
  - Draw a region r from $Multi(\pi_z)$
  - For each group member u in the member set A_{g_i}, draw u from $Multi(\psi_z)$
  - Draw a venue index v from a hybrid model $Multi(\phi_r) \times Multi(\rho_z)$
  - Draw a venue coordinate l_v from Gaussian Distribution $\mathcal{N}(\mu_r, \Sigma_r)$

- For each group-item pairs (i,j), draw the implicit response $X_{ij} \sim \mathcal{N}(G_i V_j^T, w_{ij}^{-1})$
- For each group-user social pairs (i,m), draw the implicit response $S_{im} \sim \mathcal{N}(G_i U_m^T, d_{ij}^{-1})$

4.3 Learning for GG For the group geographical model (GG), the log likelihood of the event corpus is:

(4.1)

$$\mathcal{L}_{GG} = \int p(z|\theta)p(\theta|\alpha)d\theta \int p(r|z,\pi)p(\pi|\eta)d\pi$$
$$\int p(u|z,\psi)p(\psi|\beta)d\psi \int\int p(v|r,z,\phi,\rho)p(\rho|\sigma)p(\phi|r)d\rho d\phi$$
$$p(l|r,\mu_r,\Sigma_r)$$

Gibbs sampling [5] is used to learn the known parameters in GG. We calculate the posterior probability as follows:

$$p(z_{g_i} = z, r_{g_i} = r|\mathbf{z}^{\neg g_i}, \mathbf{r}^{\neg g_i}, v, l, A_{g_i}) = \frac{n_{gz}^{\neg g_i} + \alpha}{\sum_{z'}(n_{gz'}^{\neg g_i} + \alpha)}$$
$$\times \frac{n_{zr}^{\neg g_i} + \eta}{\sum_{r'}(n_{zr'}^{\neg g_i} + \eta)} \times \prod_{u_m \in A} \frac{n_{zu}^{\neg g_i u_m} + \beta}{\sum_{u'}(n_{zu'}^{\neg g_i u_m} + \beta)} \times$$
$$\frac{n_{zv}^{\neg g_i} + \gamma}{\sum_{v'}(n_{zv'}^{\neg g_i} + \gamma)} \times \frac{n_{rv}^{\neg g_i} + \omega}{\sum_{v'}(n_{rv'}^{\neg g_i} + \omega)} \times p(l|\mu_r, \mathbf{\Sigma}_r),$$

where n_{gz} is the number of times that a topic z has been sampled from a group g and n_{zu} is the number of times that a group member u has been sampled from the distribution ψ_z specific to topic z. n_{zv} is the number of times that the event venue v has been sampled from ϕ_z specific to topic z. n_{zr} is the number of times that a region r has been sampled from the distribution π_z specific to topic z and n_{rv} is the number of times that a venue v has been sampled from ρ_r specific to r. Superscript $\neg g_i$ denotes a quantity excluding the current instance g_i.

The geographical probability of a venue coordinate l_v specific to a region r is characterized by a Gaussian distribution. The geographical Gaussian distribution $p(l|\mu_r, \mathbf{\Sigma}_r)$ is defined as follows:

$$p(l|\mu_r, \mathbf{\Sigma}_r) = \frac{1}{2\pi\sqrt{\mathbf{\Sigma}_r}} exp(\frac{-(l-\mu_r)^T \mathbf{\Sigma}_r^{-1}(l-\mu_r))}{2}), \quad (4.2)$$

where $\mu_r, \mathbf{\Sigma}_r$ are the mean and covariance vectors for a region r, respectively.

After sufficient sampling iterations, the update rules for $\{\theta, \psi, \pi, \phi, \rho\}$ are as follows:

$$\hat{\theta}_{gz} = \frac{n_{gz} + \alpha_z}{\sum_{z'} n_{gz} + \alpha_z} \qquad \hat{\psi}_{zu} = \frac{n_{zu} + \beta_u}{\sum_{u'} n_{zu} + \beta_u}$$
$$\hat{\pi}_{zr} = \frac{n_{zr} + \eta_r}{\sum_{r'} n_{zr} + \eta_r} \qquad \hat{\phi}_{zv} = \frac{n_{zv} + \gamma_v}{\sum_{v'} n_{zv} + \gamma_v} \quad (4.3)$$
$$\hat{\rho}_{rv} = \frac{n_{rv} + \omega_v}{\sum_{v'} n_{rv} + \omega_v}$$

We employ the method of moments [25] to update Gaussian distribution parameters μ_r and $\mathbf{\Sigma}_r$. The update rules for μ_r and $\mathbf{\Sigma}_r$ are defined as follows:

$$\mu_r = \frac{1}{|L_v|}\sum_{v \in L_v} l_v, \quad \mathbf{\Sigma}_r = \frac{1}{|L_v - 1|}\sum_{v \in L_v}(l_v - \mu_r)(l_v - \mu_r)^T$$

4.4 Learning for HBGG HBGG consists of two parts: the group geographical model (GG) and social-based collaborative filtering (SOCF). In HBGG, we want to learn the latent factors $\{G, U, V\}$ and parameters $\{\theta, \psi, \pi, \phi, \rho, \mu, \Sigma\}$ from the two parts. The objective of HBGG is to minimize the function: $\mathcal{L} = \mathcal{L}_{SOCF} - \mathcal{L}_{GG}$. The first term is the rating error of the SOCF part as defined in Equation 4.4, where λ_g, λ_v, λ_u and λ_s are regularized parameters. The second term is the log likelihood of the event corpus in GG defined in Equation 4.1.

(4.4)

$$\mathcal{L}_{SOCF} = \sum_{ij}\frac{w_{ij}}{2}(x_{ij} - G_i V_j^T)^2 + \frac{\lambda_g}{2}\sum_i (G_i - \theta_i)(G_i - \theta_i)^T$$
$$+ \frac{\lambda_v}{2}\sum_j V_j V_j^T + \frac{\lambda_s}{2}\sum_{im}\frac{d_{im}}{2}(s_{im} - G_i U_m^T)^2 + \frac{\lambda_u}{2}\sum_m U_m U_m^T$$

We adopt a hybrid inference procedure, combining sampling and variance optimization. We use a Gibbs-EM [21] to alternate between collapsed Gibbs sampling [5] and gradient decent for estimating parameters in our model. In the E-step, Gibbs sampling is used to learn the hidden variables $\{z, r\}$ by fixing the latent parameters $\{G, U, V\}$. In the M-step, we perform gradient decent to learn latent factors by fixing the topic and region assignments.

4.4.1 E-Step In the E-step, we apply the same learning process for the latent variables $\{z, r\}$, as the Gibbs Sampling learning process of GG to infer $\{\theta, \psi, \pi, \phi, \rho\}$ (Section 4.3). We fix the latent factors $\{G, V, U\}$ to be updated in the gradient descent step. $\{\theta, \psi, \pi, \phi, \rho\}$ are updated as in Equations 4.3.

4.4.2 M-step we perform gradient descent to learn the latent factors $\{G, V, U\}$, given the current estimate of θ in the previous E-step. The optimization procedure is similar to one-class matrix fatorization by setting the derivative to zero. The update formulae are as follows:

$$G_i \rightarrow (\lambda_s U D_i U^T + V W_j V^T + \lambda_g I_K)^{-1}$$
$$(V W_i X_i + \lambda_s U D_i S_i + \lambda_v \theta_i)$$
$$V_j \rightarrow (G W_j G^T + \lambda_v I_K)^{-1} G W_j X_j$$
$$U_m \rightarrow (\lambda_s G D_m G^T + \lambda_u I_K)^{-1}(\lambda_s G D_m S_m)$$

where W_i and D_i are diagonal matrices for a group g_i. W_i has the elements of w_{ij} on the diagonal, for $j = 1, 2, ...V$. D_i has d_{im} as its diagonal elements, for $m = 1, 2, ...U$. For each item, V_j, X_j, and W_j are similarly defined. For each user, U_m, D_m, and S_m are similarly defined.

4.5 Recommendation After learning the parameters, we can estimate the recommendation scores for a given group g_i. The recommendation scores in HBGG depend on two parts: (i) the group generative preferences inferred from the group geographic topic model (GG) and (ii) the preferences learned by social-based collaborative filtering (SOCF).

The inferred group preferences from GG are estimated by the probability score of a venue v given a group g_i as follows:

$$p(v, l_v|\theta_{g_i}, A_{g_i}) \propto \sum_z \Big[\hat{\theta}_{g_i,z}(\prod_{u \in A_{g_i}} \hat{\psi}_{zu})^{\frac{1}{|A_{g_i}|}} \hat{\phi}_{g_{zv}} \sum_r \hat{\pi}_{zr} \hat{\rho}_{rv} p(l|\mu_r, \mathbf{\Sigma}_r) \Big],$$

where $\{\hat{\theta}, \hat{\psi}, \hat{\pi}, \hat{\phi}, \hat{\rho}\}$ are defined as in Equation 4.3 and $p(l|\mu_r, \mathbf{\Sigma}_r)$ is defined as in Equation 4.2. This part can also estimate group preferences towards venues for newly formed groups (cold-start groups), by using Gibbs sampling on the given group members.

The second part (SOCF) leverages knowledge from groups in social connection. First, the estimated rating score of a venue v for a group g_i is $r^*_{ij} = G^*_i {V^*_j}^T$, where G^*_i is $\hat{\theta}_{g_i}$ for newly formed groups and the asterisk indicates the learned optimal parameters. Then, we can use social connections in location-based social networks to collaboratively learn the target group's preferences, based on the assumption that existing groups who have more social connections with the members of the target group affect the preferences of the new group. The learned S^*_{mi} is the inner product of user latent factor U^*_m and group latent factor ${G^*_i}^T$ (i.e., $S^*_{mi} = \sum_{u_m \in A_{g_i}} U^*_m {G^*_i}^T$), which represents social impact of the group g_i for the user u_m. S^*_{im} is used as social weights for aggregating the previous groups' preferences who have social connections with user u_m to form the recommendation scores r^s_{ij} in social-based CF:

$$r^s_{ij} = \frac{1}{|A_{g_i}|} \sum_{u_m \in A_{g_i}} S^*_{mi} r^*_{ij} = \frac{1}{|A_{g_i}|} \sum_{u_m \in A_{g_i}} S^*_{mi} G^*_i {V^*_j}^T$$

Finally, the integrated recommendation score of a venue v for a query group g_i is a linear combination of the two parts: $\tilde{r}_{g_i,v} = \lambda \tilde{p}(v, l_v|\theta_{g_i}, A_{g_i}) + (1-\lambda) * \tilde{r^s}_{ij}$, where $\tilde{p}, \tilde{r^s}$ indicate the normalized values of p and r^s, respectively, and λ is the parameter to control the relative weight of the two parts. As we will see in the experimental section, we use a validation test set to tune λ. The optimal setting of λ for both datasets is indicated in the parameter setting part, just after the list of compared methods. Our reported experiment results are based on the optimal values of λ.

5 Experiments

5.1 Dataset We used two real datasets for experimental evaluation, respectively Plancast [10] and Foursquare [4]. Plancast is an event-based social network. An event consists of a group of participants and the event venue with its geographical coordinates. Plancast also includes a social network in which a user can follow another user. We treat an event as a group decision, the event participants as the group members and the event venue as the group's selected venue. Foursquare is a location-based social network, where users can record their footprints; i.e., users check-in venues, giving their geographical coordinates. There is no explicit group information in Foursquare. In order to alleviate this problem, similar to previous work [16], we regard a set of friends who check-in at the same venue within one hour time-difference as a group check-in. Then, the set of friends become the group members and the checked-in venue is their selected item. For example, assume that three users u_1, u_2, u_3 are mutual friends (i.e., they form a clique in the social graph) and they checked-in a venue v within a 1-hour time difference. Then, a group event is formed in which u_1, u_2, u_3 are the group members and v is the venue.

After collecting the group events, we clean the two datasets by removing groups who have only a single member (since these correspond to venue selections by individuals) and venues, which have been visited only once. The statistics for both datasets about the number of groups, the number of users and the number of venues, after the cleaning process, are shown in Table 2. We can see that the average number of visited venues for a group in Plancast and Foursquare is very small, indicating that existing groups have fewer check-in records and most groups are cold-start groups. For both datasets, 15% of group events have been randomly marked off as the test set and 5% of group events have been used for learning the optimal values of the model parameters. The remaining 80% of the events are used as training data.

Table 2: Statistics (after preprocessing)

Dataset	Plancast	Foursquare
# of groups $\|G\|$	28077	6008
# of users $\|U\|$	38184	4150
# of venues $\|V\|$	8574	952
average # of friends for a user	40.26	10.73
average group size	9.08	2.12
average # of a group's check-in venues	1.06	1.24

5.2 Experiment Metrics We used two popular metrics for evaluating the quality of our model and its competitors, namely Precision@N and Recall@N [11, 26, 2]. Precision@N is the number of correctly predicted locations divided by the total number N of recommendations made. Recall@N is the ratio of recovered group events in the test set. Let R_g be the set of top-N recommendation items and T_g is the test data for group g. Let $\mathcal{G}_{test}$ be the set of test groups. We denote by $H_g = |R_g \cap T_g|$ the number of correctly predicted locations with regard to group g (i.e., the number of successfully predicted locations). Then, Precision@N and Recall@N are defined as:

$$Precision@N = \frac{\sum_{g \in \mathcal{G}_{test}} H_g}{N \cdot |\mathcal{G}_{test}|}, \quad Recall@N = \frac{\sum_{g \in \mathcal{G}_{test}} H_g}{\sum_{g \in \mathcal{G}_{test}} T_g}.$$

5.3 Compared Methods The first class of compared methods are score-based aggregation methods. For the methods in this class, a user-based collaborative filtering method [1] is used to estimate individual preferences. We include in our comparison three representative approaches in this class: AVE-CF, LM-CF and RD-CF [2]. AVE-CF is to average the item's preference scores by all group members. LM-CF takes the smallest score given to the item by any group member (least misery). In RD-CF, the group recommendation score is estimated by combining *relevance* and *disagreement* in the group. Relevance of the item to the group is based on AVE-CF or LM-CF, while disagreement is either the average pairwise difference of recommendation scores by group members (PD), or the variance of members' recommendation scores (VD). In experiments, we show the best performance by the different variations for combining relevances (AVE-CF or LM-CF) and disagreement scores (PD or VD). The second class of competitors are some advanced models for group recommendation, including COM [26], PIT [11] and CGAR [16]. COM and PIT are probabilistic topic models. CGAR is a hierarchical Bayesian model which combines LDA-like topic modeling with CF for group-activity recommendation. The last class of competitors are our proposed methods HBGG and GG. HBGG is the hierarchical Bayesian geographical model introduced and represented in Section 4.2. GG is the group geographical topic model in Section 4.1, a simple version of HBGG.

We used 5% of group events to learn the optimal parameter settings. For PIT, this process gives $\alpha = 50/K$, $\beta = 0.01$ and $\gamma = 0.01$. For COM, we get the following hyperparameter values: $\alpha = 50/K, \beta = \eta = \rho = 0.01$. We do the same to learn the parameters of one-class collaborative filtering (cf. Section 3.2) in CGAR and HBGG and get $c_{ij} = 1$ for positive observations and $c_{ij} = 0.01$ for negative observations. For CGAR, $\lambda_G = 1$, $\lambda_A = 0.01$ for the regularized parameters in collaborative filtering. $\alpha = \gamma = 50/K$, $\beta = \delta = 0.01$ for topic models. For our HBGG model, we get the optimal values $\lambda_g = 1$, $\lambda_u = 0.01$, $\lambda_v = 0.01$ and $\lambda_s = 0.01$ for social-based collaborative filtering. In addition, we get $\alpha = 50/K, \eta = 10/R, \beta = \gamma = \omega = 0.01$. The optimal λ value for controlling the group geographical topic model and social-based collaborative filtering is $\lambda = 0.8$ for Plancast and $\lambda = 0.7$ for Foursquare.

5.4 Experimental Results Figure 3 shows the Precision@N and Recall@N of the recommendation results on Foursquare and Plancast when K=50 and R=50 (similar results can be observed for larger K and R). We can see that our method (HBGG) clearly outperforms all its competitors. In Plancast, the improvements over the best competitor COM [26] are 28.49% in Precision@5 and 15.25% in Recall@5. The largest improvements are 47.12% in Precision@20 and 25.98% in Recall@25. In Foursquare, the improvements at Precision@5 and Recall@5 are 33.45% and 27.07%, respectively. The largest improvements are 45.66% in Precision@10 and 27.07% in Recall@5. Among the competitors, the generative models COM and PIT are better than aggregation-based approaches like CF-AVE, CF-LM and CF-RD. PIT's aggregation of group member preferences based on personal impacts is inferior to the approach followed by COM, although it brings significant improvements over CF methods. This indicates that group decisions are different from user personal preferences, and that aggregation of preferences by individual group members is insufficient. On the other hand, following a generative model like COM and our GG for the group decision process has superior performance. Our method GG outperforms COM because we consider group member engagement, group preferences, and group mobility patterns at the same time. Integrating the social-based collaborative filtering with GG (i.e., our proposed HBGG model) further improves performance. This indicates that the collaborative knowledge from social connections enhances the learning of group preferences.

5.5 Cold-Start problem Cold-Start groups which have no or few records in training set are very common, especially due to the numerous possible combinations of users that can potentially form groups. In our data, about 90.4% of the groups in the test set have no records in the training set in Plancast and about 76.25% of the groups in the test set are cold-start groups. Tables 3 and 4 report the Precision@5 and Recall@5 performance of all competitors on Cold-Start groups in both datasets (%). The last column shows the improvement of our proposed method (HBGG) over the best method from previous work. We can see that HBGG achieves higher improvements on cold-start groups in Precision@5 and Recall@5 for both datasets, when compared to the results on all groups (Figure 3). As almost 90.4% groups of the test set in Plancast are cold-start, we also see that the improvements on cold-start groups are larger than improvements on all groups, respectively (33.01% vs. 28.49% in Precision@5 and 18.92% vs. 15.25% in Precision@5). In Foursquare, the ratio of cold-start groups over test groups is relatively small and we can see that our methods have higher improvements on cold-start groups, respectively (37.08% vs. 33.45% in Precision@5 and 32.73% vs. 27.07% in Precision@5). In summary, our model is much more effective in handling cold-start groups in group recommendation, compared to the state-of-the-art methods.

5.6 Performance for Different Sizes of Groups We studied the recommendation results for groups of different sizes. As most of groups in Foursquare (more than 80%) have only two group members, we only show the recommendation performance for groups of different sizes for Plancast in Fig-

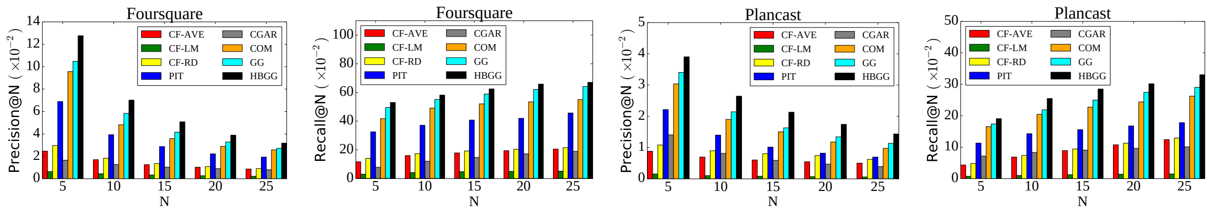

Figure 3: Precision@N and Recall@N for Foursquare and Plancast

Table 3: Precision@5 performance for Cold-Start Groups (%).

Dataset	CF-AVE	CF-LM	CF-RD	PIT	COM	CGAR	GG	HBGG	Improvement
Plancast	0.81	0.13	0.88	2.11	2.90	1.49	3.26	**3.86**	33.01
Foursquare	1.28	0.23	1.47	5.91	8.00	1.65	9.02	**10.97**	37.08

Table 4: Recall@5 performance for Cold-Start Groups (%).

Dataset	CF-AVE	CF-LM	CF-RD	PIT	COM	CGAR	GG	HBGG	Improvement
Plancast	4.30	1.31	4.38	11.62	15.46	7.43	17.08	**18.86**	18.92
Foursquare	8.21	1.46	9.34	37.78	45.16	8.03	51.19	**59.95**	32.73

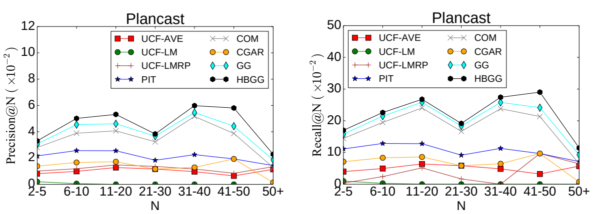

Figure 4: Performance for Different Size of Groups

Table 5: Recall@5 of HBGG varying R and K (Plancast)

R \ K	10	20	30	40	50	100
10	0.1379	0.1412	0.1512	0.1620	0.1641	0.1642
20	0.1477	0.1558	0.1619	0.1700	0.1709	0.1710
30	0.1495	0.1578	0.1636	0.1737	0.1747	0.1750
40	0.1511	0.1599	0.1737	0.1751	0.1841	0.1841
50	0.1522	0.1605	0.1768	0.1803	0.1907	0.1907
100	0.1523	0.1605	0.1768	0.1803	0.1907	0.1907

ure 4. We can see that our methods outperform all other methods for groups of different sizes. In addition, our methods have even higher improvements for groups with larger sizes. This may indicate our method can better handle large groups when making group recommendation. On the other hand, CF-based aggregation methods have poor performance when the group size is large. This happens because that aggregation strategies receive "noise" from group members of low impact within their large group.

Table 6: Recall@5 of HBGG varying R and K (Foursquare)

R \ K	10	20	30	40	50	100
10	0.3652	0.4434	0.4906	0.4953	0.5081	0.5082
20	0.3774	0.4488	0.4977	0.5085	0.5135	0.5135
30	0.3796	0.4568	0.5033	0.5102	0.5209	0.5209
40	0.3825	0.4600	0.5124	0.5198	0.5246	0.5246
50	0.3915	0.4620	0.5187	0.5254	0.5302	0.5303
100	0.3916	0.4620	0.5187	0.5254	0.5303	0.5303

5.7 Impact of Model Parameters We study the impact of model parameters on recommendation results. We vary the number of topics K and the number of regions R, and report Precision@5 and Recall@5 results of HBGG for Plancast and Foursquare in Table 5 and Table 6. The recommendation performance of HBGG increases as the number of topics K increases. When the number of topics is larger than 50, the change in the improvement is very small. Similarly, when the number of regions R is larger than 50, the results do not change much. Therefore, the results shown in Sections 5.4 and 5.5 use the parameter setting $K = 50$ and $R = 50$.

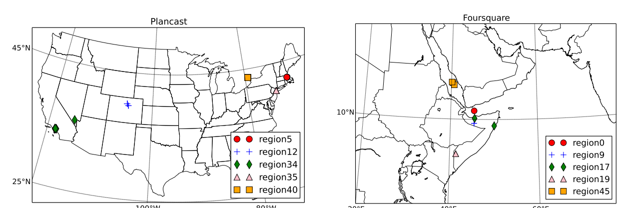

Figure 5: Geographical regions in Plancast and Foursquare

5.8 Geographical Region Analysis We also investigated the geographical regions discovered in our models (when K=50 and R=50). For both datasets, we randomly select 5 regions and plot the top-10 venues (ranked on the probability of region-venue distribution ϕ) of each region on the map, as shown in Figure 5. The regions are estimated based on group's history check-in records. We can see that the top-10 venues in each region are very close to each other geographically. This indicates that our models can cluster the venues into geographical regions and that groups tend to select some nearby venues for their activities.

6 Conclusions

In this paper, we propose a generative group geographical topic model (GG) based on group membership and group mobility regions. Experimental results show that GG outperforms some generative models which consider only topic-dependent group preferences or model individual preferences for aggregation. This indicates that the group geographical regions have important effects on inference of group preferences. Furthermore, we design a hierarchical Bayesian graphical model (HBGG) that combines the group geographical model (GG) with social-based collaborative filtering (SOCF). SOCF integrates social influences between users and existing groups into one-class collaborative filtering to enhance the learning of group latent features. The superior recommendation performance of our methods, especially for cold-start groups, indicates that the social structure is important for group recommendation and can help to alleviate cold-start issues. In the future, we plan to crawl content information for venues, e.g. category information and leverage semantic information to further alleviate the cold-start problem in group recommendation. Also, we plan to consider exploring temporal impacts on group decisions and using temporal features for dynamic group recommendation in order to further improve recommendation effectiveness.

7 Acknowledgments

This work was supported by grant GRF 17205015 and GRF 17201414 from Hong Kong RGC.

References

[1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.*, 17(6):734–749, 2005.

[2] S. Amer-Yahia, S. B. Roy, A. Chawla, G. Das, and C. Yu. Group recommendation: Semantics and efficiency. *PVLDB*, pages 754–765, 2009.

[3] L. Baltrunas, T. Makcinskas, and F. Ricci. Group recommendations with rank aggregation and collaborative filtering. In *RecSys*, pages 119–126, 2010.

[4] H. Gao, J. Tang, and H. Liu. gscorr: modeling geo-social correlations for new check-ins on location-based social networks. In *CIKM*, pages 1582–1586, 2012.

[5] T. L. Griffiths and M. Steyvers. Finding scientific topics. *PNAS*, pages 5228–5235, 2005.

[6] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM*, pages 263–272, 2008.

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

[8] H. Li, D. Wu, and N. Mamoulis. A revisit to social network-based recommender systems. In *SIGIR*, pages 1239–1242, 2014.

[9] H. Li, D. Wu, W. Tang, and N. Mamoulis. Overlapping community regularization for rating prediction in social recommender systems. In *RecSys*, pages 27–34, 2015.

[10] X. Liu, Q. He, Y. Tian, W. Lee, J. McPherson, and J. Han. Event-based social networks: linking the online and offline social worlds. In *KDD*, pages 1032–1040, 2012.

[11] X. Liu, Y. Tian, M. Ye, and W. Lee. Exploring personal impact for group recommendation. In *CIKM*, pages 674–683, 2012.

[12] Y. Liu, M. Ester, B. Hu, and D. W. Cheung. Spatio-temporal topic models for check-in data. In *2015 IEEE International Conference on Data Mining*, pages 889–894, 2015.

[13] Y. Liu, B. Zhou, F. Chen, and D. W. Cheung. Graph topic scan statistic for spatial event detection. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, CIKM '16, 2016.

[14] Z. Lu, H. Wang, N. Mamoulis, W. Tu, and D. W. Cheung. Personalized location recommendation by aggregating multiple recommenders in diversity. In *Proceedings of the Workshop on Location-Aware Recommendations, LocalRec*, 2015.

[15] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. M. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *ICDM*, pages 502–511, 2008.

[16] S. Purushotham, C. J. Kuo, J. Shahabdeen, and L. Nachman. Collaborative group-activity recommendation in location-based social networks. In *GEOGROWD*, pages 8–15, 2014.

[17] S. Purushotham and Y. Liu. Collaborative topic regression with social matrix factorization for recommendation systems. In *ICML*, 2012.

[18] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: bayesian personalized ranking from implicit feedback. In *UAI*, pages 452–461, 2009.

[19] R. Salakhutdinov and A. Mnih. Probabilistic matrix factorization. In *NIPS*, pages 1257–1264, 2007.

[20] B. M. Sarwar, G. Karypis, J. A. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *WWW*, pages 285–295, 2001.

[21] H. M. Wallach. Topic modeling: beyond bag-of-words. In *ICML*, volume 148, pages 977–984, 2006.

[22] C. Wang and D. M. Blei. Collaborative topic modeling for recommending scientific articles. In *KDD*, pages 448–456, 2011.

[23] M. Ye, X. Liu, and W. Lee. Exploring social influence for recommendation: a generative model approach. In *SIGIR*, pages 671–680, 2012.

[24] H. Yin, Y. Sun, B. Cui, Z. Hu, and L. Chen. LCARS: a location-content-aware recommender system. In *KDD*, pages 221–229, 2013.

[25] H. Yin, X. Zhou, Y. Shao, H. Wang, and S. W. Sadiq. Joint modeling of user check-in behaviors for point-of-interest recommendation. In *CIKM*, pages 1631–1640, 2015.

[26] Q. Yuan, G. Cong, and C. Lin. COM: a generative model for group recommendation. In *KDD*, pages 163–172, 2014.