Studying User Influence in Personalized Group Recommenders in Location Based Social Networks

Sanjay Purushotham
University of Southern California
Los Angeles, CA 90089
spurusho@usc.edu

C.-C. Jay Kuo
University of Southern California
Los Angeles, CA 90089
cckuo@sipi.usc.edu

Abstract

Location-based social networks (LBSNs) such as Foursquare, Google+ Local have become ubiquitous platform for users to share their activities with family and friends. LBSNs provide rich user interactions and location information to personalize recommendation systems for users, groups and places. In this paper, we propose novel collaborative filtering based hierarchical Bayesian models which jointly learns activities and group preferences to perform personalized group recommendations. We study how explicit modeling of user influence affects the personalized group recommendations. We show that well-designed personalized generative group recommenders can capture user expertise and group dynamics accurately and they out-perform state-of-the-art group recommendation systems.

1 Introduction

Location-Based Social Networks (LBSNs) such as Google+ Local, Foursquare, Gowalla[1], have become ubiquitous platform for users to share meetups/activities with family and friends. Personalized user recommendation in LBSNs has been recently studied in the past few years [9], [3], [11] to recommend activities or friends. However, many users use LBSNs to co-ordinate meet-ups/group activities and thus, there is a need for personalized group recommenders for LBSNs. Recommendation to groups (for example, recommending a movie for friends to watch together) is a challenging problem since the users of the group may or may not share similar tastes, and user preferences may change due to other users in the group. Therefore, it is important for recommendation systems to capture group dynamics such as user interactions, user group membership, user influence etc. for personalizing recommendations to groups. Personalized Group recommendation is possible for LBSNs since they provide rich content (location, activities, time-stamps) and social network information which help in accurate modeling of group dynamics.

In this paper, we propose a class of Collaborative filtering based Hierarchical Bayesian models for personalized group recommendation, and study the user influence on group recommendation. Our contributions include: (1) proposing generative probabilistic modeling framework for capturing group dynamics in group recommendation, (2) Studying the impact of model parameters and user influence on personalized group recommendation, and (3) Handling data sparsity and cold-start recommendation challenges associated with personalized group recommendation.

There is limited previous work on personalized group recommendation in LBSN. [10] proposed a probabilistic framework to model the generative process of group activities, though the final group recommendation is made by aggregation of user selections without directly learning group preferences. In [7], we proposed a probabilistic approach where we model groups and activities as generative processes to capture user-group interactions and location semantics respectively, and finally used collaborative filtering to perform group-activity recommendation. However, [7] does not

[1] Gowalla was a LBSN operational till 2012 when it was acquired by Facebook

explicitly incorporate user influence into the model which we believe is very useful for personalization of group recommendation. Motivated by our previous work, we propose a class of probabilistic models to incorporate group dynamics such as user influence, user-interactions, user-group membership into the probabilistic framework and study their effect on personalized group recommendation in LBSN. In the following sections, we describe our models, algorithm, and report our preliminary results on a real-world large LBSN (Gowalla) dataset.

2 Personalized Collaborative Group Recommender (PCGR)

In this paper, we first discuss PCGR [7] and then propose two extensions namely PCGR-D, PCGR-EF. Figure 1 shows our proposed probabilistic models. Our models belong to a class of generalized hierarchical Bayesian models and they jointly learn the group and activity latent spaces. We use topic models based on Latent Dirichlet Allocation (LDA) [2] to model the location-activity descriptions and the user-group memberships, and we use matrix factorization to match the group latent features to the latent features of locations. Our modeling framework fuses topic models with matrix factorization to obtain a consistent and compact latent feature representation.

PCGR

PCGR combines topic models of group and location activity in a collaborative filtering framework. That is, PCGR represents groups with activity interests and assumes that locations are generated by a topic model which describe activities that could be performed at a location. Furthermore, PCGR assumes that groups are generated from users who are drawn from latent communities. Thus, a user is a member of different communities, while a group is formed by multiple users belonging to different communities. In other words, a group is generated as a mixture of users from different communities. Communities are latent variables in our model and are analogous to topics in LDA, and they represent collection of users with common interests. Group generative process of PCGR captures the user-user interactions and user-group membership dynamics. PCGR additionally includes latent variables ϵ_i and ξ_j which act as offsets for community topic proportions ϕ_i and activity topic proportions θ_j respectively, when modeling the group-activity ratings. As more groups check-in at a location, we have a better idea of what these offsets are. The offset variable ξ_j can explain, for example, that an activity at location L is more interesting to a group G_1 than it is to a different group G_2. How much of the group-activity prediction relies on location, and how much it relies on other groups depend on the group preferences and how many groups have checked-in at that location. Note, we use the terms ‘location’ and ‘place’ interchangeably throughout the paper.

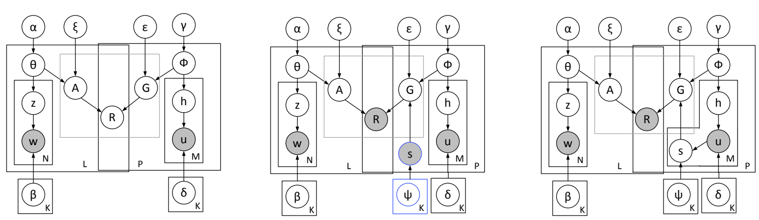

(a) PCGR (b) PCGR-D (c) PCGR-EF

Figure 1: Personalized Collaborative Group Recommendation Systems (PCGR)

Let $G \in \mathbb{R}^{K\times P}$ and $A \in \mathbb{R}^{K\times L}$ be the latent group and location-activity feature matrices in figure 1a, with column vectors G_i and A_j representing the group-specific and location-activity specific latent feature vectors respectively. The conditional distribution over the observed group location activity ratings $R^{P\times L}$ can be shown as

$$P(R|G, A, \sigma_R^2) = \prod_{i=1}^{P}\prod_{j=1}^{L} \mathcal{N}(r_{ij}|g(G_i^T A_j), \sigma_R^2)^{I_{ij}^R} \quad (1)$$

where $\mathcal{N}(x|\mu, \sigma^2)$ is the pdf of Gaussian distribution with mean μ and variance σ_R^2, and I_{ij}^R is the indicator function that is ‘1’ if group ‘i’ has rated location ‘j’, and equal to 0 otherwise. Note, here we consider group checkin as implicit group rating. The function $g(x)$ is the logistic function $g(x) = \frac{1}{1+exp(-x)}$, which bounds the range of $G_i^T A_j$ within $[0, 1]$. The group and location-activity latent feature vectors are generated as follows:

$$P(G|\sigma_G^2) \sim \mathcal{N}(\phi_i, \lambda_G^{-1} I_K) \quad (2)$$

$$P(A|\sigma_A^2) \sim \mathcal{N}(\theta_j, \lambda_A^{-1} I_K) \quad (3)$$

where $\lambda_G = \sigma_R^2/\sigma_G^2$ and $\lambda_A = \sigma_R^2/\sigma_A^2$. The generative process of PCGR model is shown below:

1. For each group i,
   (a) Draw community proportions $\phi_i \sim$ Dirichlet(γ)
   (b) Draw group latent offset $\epsilon_i \sim \mathcal{N}(0, \lambda_G^{-1} I_K)$ and set the group latent vector as $g_i = \epsilon_i + \phi_i$
   (c) For each user (u_{im}) in group i,
      i. Draw community assignment $h_{im} \sim$ Multinomial (ϕ)
      ii. Draw user $u_{im} \sim$ Multinomial($\delta_{h_{im}}$)
2. For each location j,
   (a) Draw topic proportions $\theta_j \sim$ Dirichlet(α)
   (b) Draw activity latent offset $\xi_j \sim \mathcal{N}(0, \lambda_A^{-1} I_K)$ and set the location-activity latent vector as $a_j = \xi_j + \theta_j$
   (c) For each word (w_{jn}) describing an activity at location j,
      i. Draw activity (topic) assignment $z_{jn} \sim$ Multinomial (θ)
      ii. Draw word $w_{jn} \sim$ Multinomial($\beta_{z_{jn}}$)
3. For each group-location pair *(i,j)*, draw the rating

$$r_{ij} \sim \mathcal{N}(g_i^T a_j, c_{ij}^{-1})$$

PCGR-D and PCGR-EF

In PCGR, we recommend activities based on the group preferences learned from the model. However, some users are experts (i.e. dominant users, see section 3.2 for definition) in certain activities and they have tremendous influence on their group's activities. Even though PCGR captures user influence implicitly, we could explicitly model user influence into our PCGR framework. Here, we propose two models: PCGR-D and PCGR-EF to incorporate user influence into PCGR framework. In PCGR-D, we use a deterministic switch 's_i' to select group preference according to an influential (expert/dominant) user 'i' in the group. In PCGR-EF, we use switch 's' as a latent variable to select group preference based on user influence in the group. The generative processes of PCGR-D model is similar to PCGR except for the generative process for groups. For brevity, below we only show group generative process for PCGR-D.

1. For each group i,
   (a) Draw community proportions $\phi_i \sim$ Dirichlet(γ)
   (b) Switch $s_i = \mathbb{I}(\mathcal{G}_i, \mathcal{D}_U)$,
   (c) If $s_i = 0$
      - Draw group latent offset $\epsilon_i \sim \mathcal{N}(0, \lambda_G^{-1} I_K)$
      - Set group latent vector as $g_i = \epsilon_i + \phi_i$
   (d) If $s_i = 1$
      - Draw user latent offset $\epsilon_i^{\psi} \sim \mathcal{N}(0, \lambda_U^{-1} I_K)$
      - Set group latent vector as $g_i = \epsilon_i^{\psi} + \psi_i$
   (e) For each user (u_{im}) in group i,
      i. Draw community assignment $h_{im} \sim$ Multinomial (ϕ)
      ii. Draw user $u_{im} \sim$ Multinomial($\delta_{h_{im}}$)

Where $\mathbb{I}(\mathcal{G}_i, \mathcal{D}_U)$ is an indicator function and takes value 1 if $\mathcal{G}_i \cap \mathcal{D}_U \neq \emptyset$, and takes value 0 otherwise. $\mathcal{D}_U$ is the set of dominant users in the dataset, $\mathcal{G}_i$ is the set of users in group i, ψ_i is the user-specific multinomial topic distribution and ϵ_i^{ψ} is the offset for ψ_i. Note: In PCGR-D, we consider only one dominant user for each group. In this paper, ψ_i is obtained using CTR model [8].

The generative processes of PCGR-EF model is also similar to PCGR model except for the generative process for groups. For brevity, we only present group generative process.

1. For each group i,
   (a) Draw community proportions $\phi_i \sim$ Dirichlet(γ)
   (b) For each user (u_{im}) in group i,
      i. Draw community assignment $h_{im} \sim$ Multinomial (ϕ)
      ii. Draw user $u_{im} \sim$ Multinomial($\delta_{h_{im}}$)
      iii. Draw $\lambda_s \sim$ Beta(ρ_a, ρ_b)
      iv. Draw switch $s_{u_{im}} \sim$ Bernoulli (λ_s)
      v. If $s_{u_{im}} = 0$
         - Draw group latent offset $\epsilon_i \sim \mathcal{N}(0, \lambda_G^{-1} I_K)$
         - Set group latent vector as $g_i = \epsilon_i + \phi_i$
      vi. If $s_{u_{im}} = 1$
         - Draw user latent offset $\epsilon_i^{\psi} \sim \mathcal{N}(0, \lambda_U^{-1} I_K)$
         - Set group latent vector as $g_i = \epsilon_i^{\psi} + \psi_i$

Where switch '$s_{u_{im}}$' is drawn from user-specific Bernoulli distribution with parameter λ_s and, λ_s has a Beta prior (ρ_a, ρ_b). Here, user u_{im} influences the group 'i' with probability λ_s. Throughout this paper, we only concentrate on PCGR and PCGR-D models. PCGR-EF is an ongoing work.

2.1 Parameter Learning

In this paper, we discuss the parameter learning for PCGR-D (PCGR parameter learning is similar, refer [7]). Given β, δ, ψ and s parameters, computing the full posterior of g_i, a_j, ϕ_i and θ_j is intractable. Here, we develop an EM-style algorithm to learn the Maximum-a-posteriori estimates. Maximization of the posterior is equivalent to maximizing the complete log-likelihood of $G, A, \theta_{1:L}, \phi_{1:P}$ and R given λ_G, λ_A and β, δ, ψ, s.

$$\mathcal{L} = -\frac{\lambda_G}{2}\sum_i (g_i - f_{s_i}(\phi_i,\psi_i))^T (g_i - f_{s_i}(\phi_i,\psi_i)) - \frac{\lambda_A}{2}\sum_j (a_j-\theta_j)^T(a_j-\theta_j)$$
$$+ \sum_j \sum_n log\Big(\sum_k \theta_{jk}\beta_{k,w_{jn}}\Big) - \sum_{ij}\frac{c_{ij}}{2}(r_{ij} - g_i^T a_j)^2 + \sum_i\sum_m log\Big(\sum_p \phi_{im}\delta_{p,u_{im}}\Big)$$

where $f_{s_i}(\phi_i,\psi_i) = \phi_i$ if $s_i = 0$, else $f_{s_i}(\phi_i,\psi_i) = \psi_i$; $\lambda_G = \sigma_R^2/\sigma_G^2, \lambda_A = \sigma_R^2/\sigma_A^2$ and Dirichlet priors (α and γ) are set to 1. We optimize this function by gradient ascent approach by iteratively optimizing the collaborative filtering variables g_i, a_j and topic proportions θ_j and ϕ_i. For g_i, a_j, maximization follows similar to matrix factorization [4]. Given a current estimate of θ_j, ϕ_i, taking the gradient of $\mathcal{L}$ with respect to g_i and a_j and setting it to zero helps us to find g_i, a_j in terms of $G, A, C, R, \lambda_G, \lambda_A, \psi$. Solving the corresponding equations will lead to the following update equations

$$g_i \leftarrow (AC_iA^T + \lambda_G I_K)^{-1}(AC_iR_i + \lambda_G f_{s_i}(\phi_i,\psi_i)) \tag{4}$$

$$a_j \leftarrow (GC_jG^T + \lambda_A I_K)^{-1}(GC_jR_j + \lambda_A\theta_j) \tag{5}$$

where C_i is diagonal matrix with $c_{ij}; j = 1....L$ as its diagonal elements and $R_i = (r_{ij})_{j=1}^L$ for group i. For each location j, C_j and R_j are similarly defined. Note that c_{ij} is confidence parameter for rating r_{ij}, for more details refer [6]. The equation (5) shows how activity topic proportions θ_j affects the location-activity latent vector a_j, where λ_A balances this effect. Given G and A, we can learn the activity topic proportions θ_j and community proportions ϕ_i. We define $q(z_{jn} = k) = \pi_{jnk}$ and then we separate the items that contain θ_j and apply Jensen's inequality:

$$\mathcal{L}(\theta_j) \geq -\frac{\lambda_A}{2}(a_j-\theta_j)^T(a_j-\theta_j) + \sum_n\sum_k \pi_{jnk}(\log\theta_{jk}\beta_{k,w_{jn}} - \log\pi_{jnk}) = \mathcal{L}(\theta_j,\pi_j) \tag{6}$$

The optimal π_{jnk} satisfies $\pi_{jnk} \propto \theta_{jk}\beta_{k,w_{jn}}$. Note, we cannot optimize θ_j analytically, so we use projection gradient approaches to optimize $\theta_{1:L}$ and other parameters $G, A, \pi_{1:L}$. After we estimate G,A and π, we can optimize β,

$$\beta_{kw} \propto \sum_j\sum_n \pi_{jnk} 1[w_{jn} = w] \tag{7}$$

We solve ϕ_i using variational Bayesian approach which is very similar to solving θ_j. Note, equation (7) is same as the M-step update for topics in LDA [2]

2.2 Prediction

After the optimal parameters, $G^*, A^*, \theta^*_{1:L}, \phi^*_{1:P}$ and β^*, δ^* are learned, our models (PCGR, PCGR-D) can be used for in-matrix and out-matrix prediction tasks (group-oriented recommendations). If $\mathbb{D}$ is the observed data, then both in-matrix and out-matrix predictions can be easily estimated. In-matrix prediction refers to the case where a group has not rated (visited) a location but that place has been visited by atleast one other group. On the other hand, out-matrix refers to the case where none of the groups have rated a particular place i.e. the place has no rating records (no group checkins). For in-matrix prediction, we use the point estimate of g_i, θ_j and ξ_j to approximate their expectations as:

$$\mathcal{E}[r_{ij}|\mathbb{D}] \approx \mathcal{E}[g_i|\mathbb{D}]^T(\mathcal{E}[\theta_j|\mathbb{D}] + \mathcal{E}[\xi_j|\mathbb{D}]) \tag{8}$$

$$r^*_{ij} \approx (g^*_i)^T a^*_j \tag{9}$$

In out-matrix prediction, the place is new, i.e. it does not have any ratings (cold-start recommendation setting for location). Thus, $\mathcal{E}[\xi_j] = 0$, and we predict ratings as:

$$r^*_{ij} \approx (g^*_i)^T \theta^*_j \tag{10}$$

3 Experiments

We conduct experiments to compare the performance of PCGR-D with PCGR and other state-of-the-art techniques. We evaluate our models on Gowalla dataset [1] for group-location rating prediction and activity recommendation. Our experiments help us to answer the following key questions: (a) How do our models perform with respect to the state-of-the-art group recommenders? (b) How do the model parameters (λ_G, λ_A) affect our prediction accuracy? (c) How does a dominant user (expert) influence personalized group recommendations?

3.1 Dataset Description

Our experiments were conducted on Gowalla dataset, which was a real-world LBSN. Table 1 shows the description of this dataset. Since Gowalla dataset does not contain the location content information, it was crawled from Foursquare using API and the location's geo-coordinates provided in Gowalla dataset. This dataset does not provide group checkin information. So, we infer 'groups'

Table 1: Gowalla Dataset Description

Dataset	**Gowalla**
Users	196591
Relations	950327
Check-ins	6442890
Locations	1280969
Total # of unique groups @ 1 hr	36669
# of unique groups with atleast 10 check-ins @ 1 hr	2702
# Locations checked-in by 2702 groups	94067
Total # of group checkins of 2702 groups (sparsity %)	57669 (99.9 %)

and their checkins from the user checkin information. We define a 'group' in LBSN as a set of friends checked-in at a location during a particular interval of time. Mathematically, a group in LBSN represents the 'connected components of the social network graph' during a time interval at a particular location. In [7], we did a detailed study of the Gowalla dataset and discussed Gowalla group characteristics, Group-User-Location characteristics, and User-Group membership statistics. In this paper, we first present dominant user characteristics of Gowalla dataset and then study how modeling dominant user's influence affects the personalized group recommendations.

3.2 Dominant User Characteristics

In this section, we study dominant user characteristics w.r.t group location checkins. We define 'dominant user' as a user of the group who has the most overlap in the locations visited by the group. Mathematically, we define dominant user as follows: Let X denote the set of places visited by a group G, and let Y_i denote the set of places visited by user i of the group G. Then dominant user d of G is given by $d = \arg\max_i(Y_i \cap X)$. Here, $\cap$ operation means set intersection. Let $\#(\mathcal{X})$ denote total number of elements of set $\mathcal{X}$. Let R_d denote the ratio of common places visited by a dominant user in a group w.r.t places visited by the group. Thus, R_d is given by

$$R_d = \frac{\#(Y_d \cap X)}{\#(X)}$$

The ratio R_d implies how a dominant user influences the group in terms of location checkins by the group. In our Gowalla dataset, we observe the following dominant user characteristics in terms of R_d. When $R_d \geq 0.8$, we see that 170 groups ($\sim 6\%$) have a dominant user who influences the group in (at-least) 80% of the places visited by the group. When $R_d \leq 0.50$, we see that 2000 groups ($\sim 74\%$) have a dominant user who influences the group in (at-most) 50% of the places visited by the group. When $R_d = 0$, we see that there are 80 groups ($\sim 3\%$) who have no dominant user. On the otherhand, when $R_d = 1$, we observe that 80 groups (3%) are completely influenced by the dominant user. These observations indicate that there are a few groups where dominant user has a large influence on group checkins, while there are many groups where dominant user's influence on group checkins is limited. Thus, user's influence on group checkins (and group preferences) depends on how dominant the dominant user (expert) is. In section 4.3, we study how our models utilize dominant user's influence for personalized group recommendation.

3.3 Experimental Settings

For fair comparison, our experimental settings is identical to [7]. We split Gowalla dataset into three parts - training (∼80%), held-out (∼5%) and test datasets (∼15%). The model is trained on training data, the optimal parameters obtained on the held-out data and ratings are predicted for the test data. We ran 5 simulations for all our comparison experiments. For performance evaluation, we consider (1) Average Rating Prediction Accuracy (Avg. Accuracy) and (2) Average Root Mean Squared Error (Avg. RMSE). We define the avg. accuracy for test data as the ratio of correctly predicted ratings compared to the total ratings in test data. RMSE is given by:

$$RMSE = \sqrt{\frac{1}{|\mathcal{T}|} \sum_{(i,j) \in \mathcal{T}} (\hat{R}_{ij} - R_{ij})^2}$$

Where $\hat{R}_{ij}$ is predicted ratings of group-location pairs (i, j) for a test set $\mathcal{T}$, and R_{ij} are true ratings. The average of RMSE for the test set is denoted by Avg. RMSE. Having higher avg. accuracy and lower avg. RMSE means better recommendation model.

4 Results

In this section, we present our experimental results and answer the questions raised in section 3.

4.1 Performance Comparisons

Table 2 shows the performance comparison of our models (PCGR, PCGR-D) w.r.t the state-of-the-art group recommendation systems. Both PCGR and PCGR-D outperform all the other models for in-matrix and out-matrix prediction tasks, while PCGR-D performs slightly better than PCGR. For PCGR-D, results are reported for $R_d = 1$. In section 4.3, we study impact of R_d on PCGR-D model.

Table 2: Performance Comparison on test data

	Best Aggregation method (Averaging)	**PMF** [5]	**CTR** [8]	**PGCR**	**PCGR-D**
In-matrix Avg. Accuracy	0.75	0.74	0.75	0.89	0.90
In-matrix Avg. RMSE	0.51	0.50	0.49	0.32	0.31
Out-matrix Avg. Accuracy	-	-	0.40	0.52	0.53
Out-matrix Avg. RMSE	-	-	0.77	0.69	0.68

4.2 Impact of λ_G and λ_A parameters

Our models (PCGR and PCGR-D) allow us to study how the group parameter λ_G and location parameter λ_A affects the overall performance of personalized group-activity recommendation. Figure 2 shows the impact of λ_G and λ_A on PCGR model's prediction accuracy for Gowalla test dataset. We observe that when λ_A is fixed and finite, and $\lambda_G = 0$, our model collapses to the CTR model which uses topic modeling (LDA) and group-location rating matrix for prediction. When $\lambda_G \to \infty$ (and λ_A is fixed and finite), then our model only uses group preferences for prediction without considering the activities offered at a location. When $\lambda_A = 0$ and $\lambda_G = 0$, then our model reduces to the Matrix Factorization (PMF) [5], since the latent variables G and A are not affected by the topic models. For all other cases, PCGR fuses information from activity and group topic models to predict ratings for the groups. From figure 2, we see that very small and very large values of λ_G and λ_A do not improve the prediction accuracy. This can be explained as follows: very small values (< 0.001) of parameters mean that model is close to PMF and does not use location-activity/community topics for recommendation, while very large parameter values (> 10) means that model heavily relies on the topic models and less on the collaborative filtering which is not good especially if groups/location descriptions are noisy/sparse. Our PCGR and PCGR-D models obtain best prediction accuracy when $\lambda_G \sim 0.01$ and $\lambda_A \sim 0.1$, showing that both group communities and location activity topics are important for better recommendation. We also observe that higher λ_A (compared to λ_G) gives better prediction since we have a better and reliable descriptions for location-activity topic models.

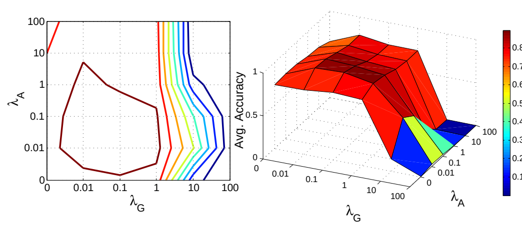

Figure 2: Impact of λ_G and λ_A on prediction accuracy for PCGR model

4.3 Dominant User Influence

Figure 3 shows how dominant user's influence affects the performance (in-matrix avg. accuracy) of PCGR-D on test data. To study the effects of explicitly incorporating user influence into the model, we select different groups in the test data by varying R_d. When $R_d = 1$ dominant user completely influences group check-ins. For $R_d = 1$, there are around 35 groups with dominant users in our test data, and PCGR-D chooses user-specific multinomial topic distribution ψ (learned from CTR [8]) for group recommendation to these groups. When $R_d = 0$, a few groups (~ 10) do not have a dominant user, and only these groups use group-specific topic distribution ϕ for group recommendation, while the remaining groups use ψ for group recommendation. We observe that PCGR-D performs slightly better than PCGR when $R_d > 0.8$ since it explicitly incorporates dominant user's influence into the recommendation model. However, as R_d decreases, the performance of PCGR-D decreases compared to fixed PCGR. This is because when $R_d < 0.5$, PCGR-D relies on (dominant) user-specific topic distributions (ψ) for group recommendation to many groups, while it uses group topic distribution for only a few groups. When $R_d = 0$, PCGR-D model's accuracy is closer to CTR model's accuracy since PCGR-D uses dominant user's latent vector ψ (instead of group latent vector) for group recommendation, but ψ does not correctly capture the group preferences (CTR does not model group dynamics). Thus, we find that explicitly modeling dominant user's influence is useful for personalized group recommendation only when $R_d (> 0.8)$ is higher (i.e. really dominant users).

4.4 Discussion

In our paper, we showed that modeling group dynamics such as user interactions, user influence & user group memberships using group generative process helps us to personalize group preferences. For the activity descriptions at a location, we only used the information provided by the LBSN; and we considered group-checkin as a indicator for group rating (implicit rating) of the location. However, it is possible that groups checkin at a location but do not like the activities offered there. Hence,

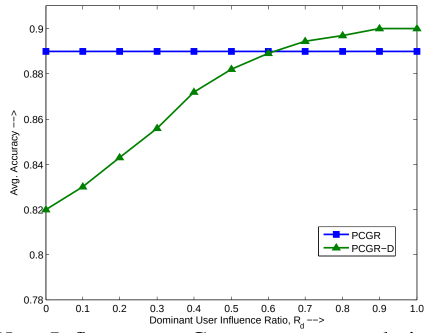

Figure 3: Dominant User Influence on Group recommendation for entire test data

we feel that mining group's feedback such as comments, actual ratings, tweets etc., provides better data to model group preferences. Also, external factors such as weather, holidays, discounts etc. could influence group activities. We could include these factors into our activity topic model using external variables, but this could make our models complex and inference harder. It is worth noting that our models provide interpretable results [7] which are useful for user studies in deployment of personalized group recommender systems.

5 Summary and Future Work

In this paper, we presented Collaborative filtering based hierarchical Bayesian models that exploit group checkins, location information, user-user interactions and user influence to learn group preferences and recommend personalized activities to groups. Our experiments on Gowalla dataset showed that our models consistently outperform the state-of-the-art group and content recommmendation systems. Our framework models the group dynamics and allows us to address cold-start recommendation for new locations/groups. The main contributions of our paper include: 1) demonstrating the effectiveness of modeling group dynamics such as user interactions, user influence, user-group membership to improve and personalize group recommendation in LBSNs, 2) studying the impact of modeling groups and locations using latent factor models.

There are many directions for our future work. First, we will study PCGR-EF model and then work on making our algorithms scalable to ever evolving LBSNs. We will conduct user studies to check subjective performance of our personalized group recommendation systems.

References

[1] http://snap.stanford.edu/data/#locnet.

[2] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 2003.

[3] B. Hu and M. Ester. Spatial topic modeling in online social media for location recommendation. In *ACM RecSys*, 2013.

[4] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 2009.

[5] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *NIPS*, 2007.

[6] S. Purushotham, Y. Liu, and C.-C. J. Kuo. Collaborative topic regression with social matrix factorization for recommendation systems. *ICML*, 2012.

[7] S. Purushotham, J. Shahabdeen, L. Nachman, and C.-C. J. Kuo. Collaborative group-activity recommendation in location-based social networks. In *ACM SIGSPATIAL, GeoCrowd*, 2014 (To appear).

[8] C. Wang and D. M. Blei. Collaborative topic modeling for recommending scientific articles. In *ACM SIGKDD*, 2011.

[9] M. Ye, P. Yin, W.-C. Lee, and D.-L. Lee. Exploiting geographical influence for collaborative point-of-interest recommendation. In *ACM SIGIR*, 2011.

[10] Q. Yuan, G. Cong, and C.-Y. Lin. Com: a generative model for group recommendation. In *ACM SIGKDD*, 2014.

[11] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang. Collaborative location and activity recommendations with gps history data. In *WWW*, 2010.