Who Wants to Join Me? Companion Recommendation in Location Based Social Networks*

Yi Liao[1], Wai Lam[1], Shoaib Jameel[2], Steven Schockaert[2], Xing Xie[3]
[1]Dept. of Systems Engineering and Engineering Management, The Chinese University of Hong Kong
[2]School of Computer Science and Informatics, Cardiff University
[3]Microsoft Research
{yliao, wlam}@se.cuhk.edu.hk
{jameels1,schockaerts1}@cardiff.ac.uk xingx@microsoft.com

ABSTRACT

We consider the problem of identifying possible companions for a user who is planning to visit a given venue. Specifically, we study the task of predicting which of the user's current friends, in a location based social network (LBSN), are most likely to be interested in joining the visit. An important underlying assumption of our model is that friendship relations can be clustered based on the kinds of interests that are shared by the friends. To identify these friendship types, we use a latent topic model, which moreover takes into account the geographic proximity of the user to the location of the proposed venue. To the best of our knowledge, our model is the first that addresses the task of recommending companions for a proposed activity. While a number of existing topic models can be adapted to make such predictions, we experimentally show that such methods are significantly outperformed by our model.

1. INTRODUCTION

Fuelled by the popularity of mobile devices, location based social networks (LBSNs) such as Foursquare have started to flourish in recent years. Users on such social networks use an app to report the venues they are visiting, typically in real time. LBSN providers thus obtain valuable data about their users in the form of lists of so-called check-in records, i.e. records of the time and venue of each reported visit. These lists can then be used to recommend venues that might be of interest to the user, and to target advertising, among others. For example, a large number of authors have looked at the problem of place-of-interest (POI) recommendation, i.e. the problem of recommending new venues based on a user's past check-in behavior [23, 22, 50, 20, 24, 51, 8, 19].

In this paper, we study a new task for LBSNs which we call companion recommendation. In particular, we consider the scenario of a user who is planning a visit (e.g. going to the movies, going to the park for a picnic, going to a concert) and who is looking for friends who would like to join. In such a case, the user may post a message such as "I'm going to the beach on Saturday. Who wants to join me?". The task we address in this paper is to predict who among a given user's friends is most likely to be interested in joining the proposed activity. The core idea is that we can cluster each user's friends based on the interests they share with the user. For example, a user may have colleagues with whom they have lunch on weekdays, friends with whom they go to concerts, other friends who share a passion for hiking, and yet other friends with whom they go out on Saturday evenings. Our proposed model will automatically induce the different kinds of friendship types that are found in a given LBSN from the past check-in behaviour of the users, and will allow us to predict which of these friendship types is most closely related to the proposed visit. This model can be used in various ways by an LBSN. For example, when a user is posting a message announcing a planned visit, the system can automatically recommend groups of friends with whom this message could be shared, as depicted in Figure 1. Along similar lines, the predictions made by the model could feed into the ranking algorithm that is used to display news feeds, i.e. the message could be given more prominence in the news feeds of friends who are more likely to be interested in joining. Finally, the model could also directly recommend companions to the user, which could be useful in some cases as users are not always aware of all interests of their online friends, given that many friendship ties on social networks are weak ties [6, 11].

Several authors have looked at the friend recommendation task

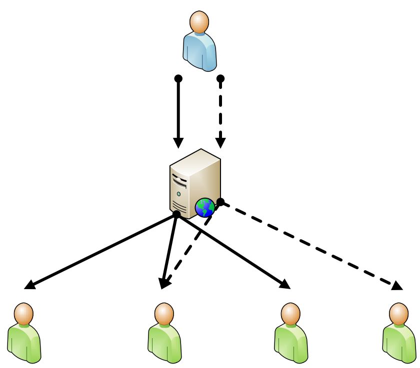

Figure 1: The solid line represents the message "I'm going to the beach on Saturday. Who wants to join me?"; The dashed line represents the message "Tonight at the movie. Who wants to come along?"

*We thank Dr. Zhiyuan Cheng, Google for providing the raw check-in dataset and Prasanta Bhattacharya of the National University of Singapore for giving us valuable information about crawling data.

ICTIR '16, September 12-16, 2016, Newark, DE, USA
© 2016 ACM. ISBN 978-1-4503-4497-5/16/09...$15.00
DOI: http://dx.doi.org/10.1145/2970398.2970420

[13, 41, 44, 34], where the aim is to predict missing friendship relations, based on information collected from social networks. While this task also involves friendship relations, it is clearly different from companion recommendation, as it does not take into account the characteristics of a specific venue, which is a key element in our proposed setting. Another related task is predicting the strength of existing friendship relations [43, 37, 16]. The framework proposed by Xiang et al. [43], for instance, makes use of the users' check-in history and previous interactions to predict friendship strength. However, such models cannot solve the companion recommendation tasks in a satisfactory way, as again they do not take into account what exactly are the shared interests between two friends. For example, two users may be close friends but enjoy very different types of music, in which case they should not be recommended to each other when planning to attend a concert. Companion recommendation is furthermore different from the task of group recommendation [35, 45, 29, 3, 33], which also involves multiple users, but where the aim is to recommend the most satisfactory venue to a group of users. Finally, the proposed companion recommendation tasks is also clearly different from POI recommendation. To the best of our knowledge, no previous works have investigated the task of recommending companions to a user planning to visit a given venue.

To tackle the companion recommendation task, we propose a probabilistic graphical model which captures the inter-relationships among the essential elements of this problem: venues, users and friendship relations. Our framework uses a first set of latent variables to model user interests and a second set of latent variables to model friendship types. User interests are modelled as distributions over venue categories, whereas friendship types are modelled as distributions over users that share particular interests. The latent variables are estimated based on the categories of the venues that have previously been visited by all users, and, for each pair of users, the list of venues which they have both visited in the past. An important underlying assumption is that the categories of the venues which two friends have both visited semantically characterizes their friendship type. Furthermore, our model also takes into account the distance between the location of the proposed venue and the location of each candidate companion (estimated from the locations of the venues which they have visited in the past). This is important, as users are clearly more likely to visit places in their vicinity. This observation, which is known as geographical mobility, has been utilized in various LBSN-based approaches [23, 55, 49, 5, 17, 42, 15, 31, 12, 54, 26].

The remainder of this paper is structured as follows. In the next section, we discuss in more detail how our model is related to existing methods. Section 3 then introduces our model, and explains how its parameters can be estimated and how it can be used to make predictions. Finally, in Section 4 we evaluate our model by comparing it to a number of baselines. While there are no existing method for the task of companion recommendation that against which we can compare directly, our evaluation will demonstrate that straightforward modifications to existing models cannot offer a competitive solution.

2. RELATED WORK

As mentioned in the introduction, we are not aware of any existing approaches that solve the task of companion recommendation. In this section, we discuss existing models for the three most closely related tasks: friendship prediction, POI recommendation and group recommendation.

2.1 Friendship Prediction

The problem of modelling how social networks evolve has attracted a lot of attention in recent years. Among others, several models have been proposed to predict which friendship relations are likely to be formed. A common approach is to treat friendship prediction as a classical link prediction problem and rely on proximity measures in the social network graph. An early example of such an approach, in the context of co-authorship networks, is presented in [25], where a model that takes into account network topology was shown to perform substantially better than random guessing. Other methods for measuring the proximity between two users include the common neighbors, Jaccard coefficient, and Admic/Adar methods, as surveyed by [25]. While the aforementioned methods look at the local neighbourhood structure, methods such as random walk approaches [2, 21] can be regarded as global proximity methods [48]. For example, Backstrom et al., [1] propose a supervised random walk algorithm for friendship prediction, incorporating users' attributes. Some studies have also found interesting patterns in existing friendship networks. For example, triadic closure is a typical structural pattern in friendship networks and has been investigated by [37, 32, 30]. Some models also focus on semantic features. In [36], for instance, the authors present a model that predicts friendship relations based the similarity between user profiles. Recent work such as [38] studies the top-k link prediction problem in social networks, which emphasizes the precision of top-k users in the recommended ranking list. Some other works such as [47, 49] study the link prediction problem by transferring information from aligned multi-networks, where multiple social networks are partially aligned at the same time. Subbian et al., [39] build a robust and effective classifier for link prediction using multiple auxiliary networks. In multi-networks, users can be extensively correlated with each other by various connections.

The aforementioned works only deal with binary friendship. However, some authors have also considered the problem of predicting the strength of existing friendship relations [16, 43]. For example, Xiang et al., [43] propose a framework to infer latent friendship strength from the similarity of the users' profiles and the frequency of their interactions. In contrast with all the existing work on friendship prediction, our model takes into account different types of friendship links to tailor companion recommendations to the characteristics of the specific venue being visited.

2.2 POI Recommendation

POI or venue recommendation is a widely studied problem in the context of LBSNs. In most models, spatial information plays a prominent role, since the probability for a user to visit a venue is closely related to the distance the user needs to travel to reach the venue (as suggested by Tobler's First Law of Geography). In [20, 53, 18, 17] the authors study GPS records, encoded as a series of time points with associated geo-coordinates, to capture patterns of user movement. Lian et al., [24] propose a framework based on matrix factorization, augmenting latent factors with vector representations, to capture the so-called activity areas of users and the influence areas of POIs. Temporal aspects of venue recommendation are studied in [9, 46], which take account of the fact that people's activities and movements vary over time.

With online social networks increasingly storing users' present and past movements, content-based venue recommendation has recently attracted much attention. In [15] the authors explore a spatial topic modeling approach to predict future venues of interest based on the textual content of a user's posts. Liu et al., [27] exploit various aspects of venue profiles and develop a joint model for venue recommendation. More recently, Gao et al., [10] propose

Symbol	Description
u, U	user and user set respectively
$\|M\|$	size of check-in records for a particular user
v, V	venue and venue set respectively
l, L	category label and category label set respectively
c, C_u	companion and companion candidates for user u
$f, \|F\|$	friendship type and number of friendship types
g	geographical location of a venue
H	historical check-in data for a companion
θ	user-specific distribution on topics
$z, \|Z\|$	topic associated with each visit and number of topics respectively
Φ^f	topic-specific distribution over friendship type
Φ^v	topic-specific distribution over venues
Φ^l	distribution over category labels specific to friendship type
$\Phi^c_{u,f}$	preference on companions specific to pairs of user u and friendship type f

Table 1: Some basic notations

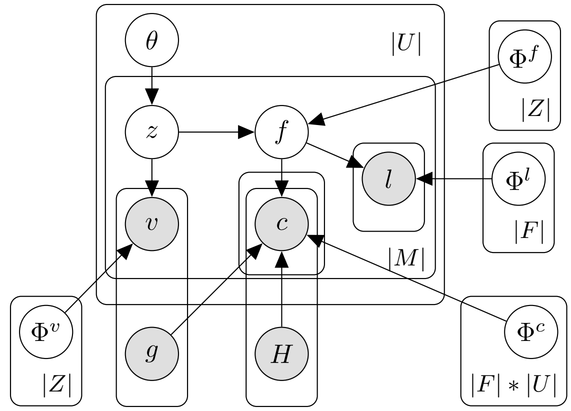

Figure 2: Graphical model for companion recommendation

a POI recommendation framework by relating three types of content information to different aspects of users' check-in behavior. Lian et al. [23] propose an Implicit-feedback based Content-aware Collaborative Filtering (ICCF) framework to incorporate semantic content and steer clear of negative sampling.

2.3 Group Recommendation

Group recommendation methods [45, 29, 3, 33, 52] aggregate the preferences of a group of users and seek to recommend venues that are most suitable for the group as a whole. For example, Yuan et al., [45] propose a generative model that studies different influences for users in a group. Cheng et al., [3] investigate multiple user behaviors in group recommendation. Salehi-Abari and Boutilier [35] have developed probabilistic inference methods for predicting individual preferences and exploit these predictions to make group decisions or recommendations based on techniques from the field of social choice theory. Given that the group recommendation task is about recommending venues, it is clear that this task is different from the problem we discuss in this paper.

3. OUR FRAMEWORK

We assume that the following data is available to support the task of companion recommendation: a list of previous check-ins for each user, a profile document for each venue, and the (generic) friendship relation between the users. We assume that each check-in record consists of (i) the venue which the user has visited, and (ii) a list of friends from the user who have also visited that venue. Furthermore, the venue profile document consists of a set of categories and a location (in the form of geo-coordinates). The set of categories in a venue profile document can contain up to three labels that describe the nature of the venue (e.g. Indian restaurant, hotel, primary school). Given the aforementioned information, and a given input user u and venue v, the task is to predict which of u's friends are most likely to be interested in joining the visit to v.

Our proposed framework is a probabilistic graphical model capturing the inter-relationships between user interests, friendship types and check-in records. In particular, we consider a generative process, in which a user u has a specific interest in mind, which is encoded as the latent topic z. We can think of z as representing the purpose of the planned visit, e.g. going to a restaurant with an Asian cuisine. The interests of a given user are sampled from a multinomial distribution, which models the overall interests of that user. Each topic z is viewed as a soft cluster of venues v, which are intuitively the venues that the user might choose to visit to satisfy z. For example, if z represents the purpose of going to a restaurant with Asian cuisine, the the associated cluster of venues will, among others, include specific Thai restaurants.

To identify candidate companions, we introduce the notion of a friendship type. Our model uses latent variables to encode these friendship types. Furthermore, each friendship type f is related to the latent topics that are used for modelling user interests. For example, there might be a friendship type for *Asian Food*, a friendship type for *Hiking*, etc. This captures the intuition that the suitable companions for a given visit depend on the characteristics of the place being visited. For example, some friends who often join the user to the gym, might not be interested in joining a visit to a Thai restaurant. Different friendship types may be compatible with a given topic to some degree. Consequently, we assign to each topic z a multinomial distribution over friendship types, capturing the suitability of different friendship types for a proposed visit.

After a suitable friendship type f is determined, we further relate the friendship type f to the list of friends from the user who have also visited that venue, denoted as companions c. Though they may not have visited the venue together with the user so far, they are potential companions when the user visits a venue next time. Friendship relations between the user and his friends, encoded as friendship type, are learned by our model, which is further used in companion recommendation in the future. For example, next time the user plans to visit a restaurant with an Asian cuisine, friends whose friendship relation with the user is identified as friendship type for *Asian Food*, is recommended as companion for this visit.

Different types of evidence are taken into account to learn the latent representations. First, we use the category information of the venues visited by each user to estimate the topics in which they are interested. Given that friendship types are related to these topics, the category information thus indirectly also serves to assign meaning to these friendship types. Second, to recommend companions for a given visit, the model takes into account the geographic location of each candidate companion, in addition to a preference score derived from the latent friendship model.

3.1 Description of the model

Figure 2 shows the graphical model, while some basic notations are explained in Table 1. Note that for clarity we do not show the hyper-priors of the parameters in the graphical model, focusing instead on the main elements of the model only. The outer rectangle is a plate representing a set of users and the inner rectangle is a plate representing the set of check-ins of a particular user. A latent topic assignment z for the check-in record is first sampled from a user-specific multinomial distribution θ, representing the user's interest. The selected topic is then used to generate both a venue and a friendship type. Specifically, the topic z is associated with a multinomial distribution Φ_z^v from which a specific venue is sampled, as well as a multinomial distribution Φ_z^f from which a specific friendship type is sampled.

To ensure that the friendship types can be characterized in terms of venue categories, we assume that for each friendship type f we have a multinomial distribution Φ_f^l from which venue categories are sampled. Note that a venue may belong to several categories. Thus a category label is sampled several times. Note that in this way, the categories of the venue being visited act as a soft constraint on the kinds of friendship types that might be selected.

To sample a companion c, we take into account both the selected friendship type and the geographic location of the user and proposed venue. The friendship type is taken into account by associating a preference vector $\Phi_{u,f}^c$ with each user and each friendship type. Note that the vector $\Phi_{u,f}^c$ is not a multinomial distribution, since the companion c is sampled with a different type of distribution in contrast with the other nodes, which will be described later. This vector captures which friends are most likely to be companions given the specific user u and the specific friendship type f. Note that we only consider companion candidates from the user's friend list, as we assume that users are only interested in visiting the venue with people they already know.

In addition to the friendship type, geographical information also plays an important role. For example, it does not seem useful to recommend a user who lives in New York as a companion when the target venue is a gym located in California. To take into account geographical mobility, for each candidate companion c, we take into account the set $\mathbf{H}_c$ containing the locations of all c's previous check-ins. Specifically, the geographical compatibility between companion c and venue v is computed by averaging the negative exponential value of the distance between the location of venue v and the venues in $\mathbf{H}_c$, as expressed by the following equation:

$$G(c,v) = \frac{1}{|\mathbf{H}_c|} \cdot \sum_{v' \in \mathbf{H}_c} exp(-||g_{v'} - g_v||^2) \qquad (1)$$

Note that, intuitively, a venue that is near the venues in $\mathbf{H}_c$ will get high compatibility score. Let $P(c|f, \mathbf{H}_c, g_v, \Phi^c)$ denote the probability that friend c is chosen as companion given the friendship type f and venue location g_v. The estimation of this probability is inspired by the Sparse Additive Generative Model (SAGE) [7, 42]. When a variable is affected by multiple facets, a common approach is to design a weighted addition of the multiple facets. In our case, the addition would be $P(c|f, \mathbf{H}_c, g_v, \Phi^c) \propto \lambda + (1-\lambda)G(c,v)$, where λ is a parameter that needs to get estimated. However, the estimation of λ could make the inference procedure complicated and sometimes may even result in overfitting. The idea of SAGE is to combine multiple facets through simple addition in log space, avoiding the inference of switching parameter. Eisenstein et al., [7] has demonstrated its applicability in many complex multifaceted generative models. Following the idea, $P(c|f, \mathbf{H}_c, g_v, \Phi^c)$ is designed as follows:

$$P(c|f, \mathbf{H}_c, g_v, \Phi^c) = \frac{exp(\Phi_{u,f,c}^c + G(c,v))}{\sum_{c' \in C_u} exp(\Phi_{u,f,c'}^c + G(c',v))} \qquad (2)$$

We can now summarize the generative process as follows.

1. For each user u, draw a distribution $\theta_u \sim$ *Dirichlet*(α) on topics, where α is a hyper-prior.
2. For each friendship type f, draw a distribution on category labels $\Phi_f^l \sim$ *Dirichlet*(κ), where κ is a hyper-prior.
3. For each topic
   (a) Draw a distribution on venues, $\Phi_z^v \sim$ *Dirichlet*(ξ),where ξ is a hyper-prior.
   (b) Draw a distribution on friendship type, $\Phi_z^f \sim$ *Dirichlet*(ψ), where ψ is a hyper-prior.
4. For each check-in record of the user u:
   (a) Draw a topic $z \sim$ *Multinomial*(θ_u)
   (b) Draw the friendship type and companion as follows:
      i. Draw a friendship type $f \sim$ *Multinomial*(Φ_z^f)
      ii. For each companion, i.e. friend of the user who have also visited the venue, draw c according to (2).
      iii. For each category label of the venue, draw $l \sim$ *Multinomial*(Φ_f^l)
   (c) Draw a check-in venue $v \sim$ *Multinomial*(Φ_z^v)

3.2 Parameter Estimation

The complete likelihood of the model is as follows:

$$\begin{aligned} &P(z, f, v, \boldsymbol{l}, \boldsymbol{c} | \alpha, \psi, \xi, \beta, \kappa, g_v, \mathbf{H}_c) \qquad (3) \\ &= \int p(z|\theta) p(\theta|\alpha) d\theta \cdot \int p(f|z, \Phi_{z,f}^f) p(\Phi_{z,f}^f) d\Phi_{z,f}^f \cdot \\ &\int p(v|z, \Phi_{z,v}^v) p(\Phi_{z,v}^v|\xi) d\Phi_{z,v}^v \cdot \\ &\prod_l^{\boldsymbol{l}} \int p(l|f, \Phi_{f,l}^l) p(\Phi_{f,l}^l|\kappa) d\Phi_{f,l}^l \cdot \prod_c^{\boldsymbol{c}} P(c|f, \mathbf{H}_c, g_v, \Phi^c) \end{aligned}$$

where $\boldsymbol{l}$ is the set of category labels l of the venue and $\boldsymbol{c}$ is the set of companions c. Note that the last factor $P(c|f, \mathbf{H}_c, g_v, \Phi^c)$ refers to (2). In contrast to the other factors, it does not integrate over Φ^c, as the companion c is not sampled from a multinomial distribution. Furthermore, due to the specific way in which we have defined $P(c|f, \mathbf{H}_c, g_v, \Phi^c)$, Gibbs Sampling [14], which has been widely used for the inference of many probabilistic graphical models, cannot to be directly applied to our proposed model. To cope with this, we apply the Gibbs EM algorithm [40] for parameter estimation.

E-step: The latent topic assignment and friendship type assignment are sampled by fixing all the parameters. To sample the topic, the standard Gibbs Sampling method for Latent Dirichlet Allocation is employed. Let $n_{u,z}$ denote the number of times the topic z is assigned to the user u. Let $n_{z,f}$ denote the number of times the friendship type f is assigned to the topic z. Let $n_{z,v}$ denote the number of times the venue v is assigned to the topic z. Let $n_{f,l}$ denote the number of the instances of the category label l that is assigned to the friendship type f. Let $|L|_v$ denote the total number of category labels of the venue and let $|C|_i$ be the number of companions for this check-in record. As usual, when we place the $\neg$ symbol, it means that the counts exclude the current case; when we place the $\cdot$ symbol, it means that the counts for all possible values of the missing index are added up. For example, $n_{f,\cdot}$ is the total number of category labels that are assigned to the friendship type

f. Topics are sampled as follows:

$$\begin{aligned} & P(z_i = j | z_{\neg i}, f, v) \\ \propto \quad & (n_{u,j}^{\neg i} + \alpha) \cdot \left(\frac{n_{j,f_i}^{\neg i} + \psi_{f_i}}{n_{j,\cdot}^{\neg i} + \psi} \right) \cdot \left(\frac{n_{j,v_i}^{\neg i} + \xi_{v_i}}{n_{j,\cdot}^{\neg i} + \xi} \right) \end{aligned} \quad (4)$$

Friendship types are sampled, given the fixed parameter Φ^c, as follows:

$$P(f_i = j | f_{\neg i}, z, u, c, l) \propto (n_{z_i,j}^{\neg i} + \psi_j) \cdot$$
$$\left(\prod_{\varsigma=1}^{|L|_v} \frac{n_{j,l_\varsigma}^{\neg i} + \kappa_{l_\varsigma}}{n_{j,\cdot}^{\neg i} + \kappa} \right) \cdot \left(\prod_{\varsigma=1}^{|C_u|_i} P(c_\varsigma | f, \mathbf{H}_{c_\varsigma}, g_v, \Phi^c) \right) \quad (5)$$

M-step: We optimize the parameter Φ^c to maximize the logarithm of the complete likelihood denoted in Equation 3. We employ the quasi-Newton method [28] to solve the problem. This is an iterative algorithm to find local maxima or minima, which has a higher computational efficiency than the standard Newton method. The gradients of the log-likelihood regarding parameter Φ^c are calculated as follows:

$$\frac{\partial L}{\partial \Phi_{u,f,c}^c} = n_{u,f,c} - \sum_{i \in M_{u,f}} \frac{exp(\Phi_{u,f,c}^c + G(c, v_i))}{\sum_{c' \in C_u} exp(\Phi_{u,f,c'}^c + G(c', v_i))} \quad (6)$$

where $n_{u,f,c}$ denotes the number of times the companion c is assigned to the pair (u, f). $M_{u,f}$ denotes the the set of check-in records that have been assigned to the friendship type f for the user u.

The remaining parameters can be estimated by using standard Gibbs sampling [14] for topic modeling. After a sufficient number of iteration, the other parameters are calculated as follows:

$$\theta_{u,z} = \frac{n_{u,z} + \alpha_z}{\sum_z^Z (n_{u,z} + \alpha_z)} \quad (7)$$

$$\Phi_{z,f}^f = \frac{n_{z,f} + \psi_f}{\sum_f^F (n_{z,f} + \psi_f)} \quad (8)$$

$$\Phi_{z,v}^v = \frac{n_{z,v} + \xi_v}{\sum_v^V (n_{z,v} + \xi_v)} \quad (9)$$

$$\Phi_{f,l}^l = \frac{n_{f,l} + \kappa_l}{\sum_l^L (n_{f,l} + \kappa_l)} \quad (10)$$

3.3 Companion Recommendation

Given a user u who wants to visit a venue v, the output of the considered task is a ranked list of recommended companions. Consider a user u who plans to visit a venue v. To generate suitable recommendations, we make use of the category information $\boldsymbol{l}$ of the venue for the inference of friendship types. The friendship type f is chosen with the following probability:

$$P(f | u, v, \boldsymbol{l}) = \sum_{z=1}^{|Z|} P(f | z, \boldsymbol{l}) P(z | u, v) \quad (11)$$

where $P(z|u, v)$ denotes the probability of drawing the topic z given the user u and the venue v and $P(f|z, \boldsymbol{l})$ denotes the probability of drawing the friendship type f given the topic z and the category labels $\boldsymbol{l}$. They can be derived using Bayes' theorem as follows:

$$P(z|u, v) \propto P(z|u)P(v|z) = \theta_{u,z}\Phi_{z,v}^v \quad (12)$$

$$P(f|z, \boldsymbol{l}) \propto P(f|z)P(\boldsymbol{l}|f) = \Phi_{z,f}^f \prod_{i=1}^{|L|_v} \Phi_{f,l_i}^l \quad (13)$$

Number of	NYC	CA	USA
Users	2219	2692	16872
Venues	5588	7828	58205
Check-in Records	54247	67689	470074
Avg. Records/User	24.45	25.14	27.86
Avg. Friends/User	26.89	29.51	40.51
Avg. Companion/Record	2.11	3.25	2.18
Avg. Category Labels/Venue	1.72	1.55	1.31

Table 2: Statistics of datasets

where $|L|_v$ is the total number of category labels of the venue.

The overall probability that the companion c is selected can then be computed as follows:

$$P(c|u, v, \boldsymbol{l}, \mathbf{H}_c, g_v, \Phi^c) = \sum_{f=1}^{|F|} P(c|f, \mathbf{H}_c, g_v, \Phi^c) P(f|u, v, \boldsymbol{l})$$

where $P(c|f, \mathbf{H}_c, g_v, \Phi^c)$ is evaluated as in (2) and $P(f|u, v, \boldsymbol{l})$ is evaluated as in (11).

4. EXPERIMENTS AND RESULTS

4.1 Experimental Setup

We conduct experiments on real-world check-in data obtained from Foursquare, a location based social network. We obtained a raw dataset[1] from the first author of [4]. It contains a tuple (UserID, VenueID, Location) for each check-in record, which was collected from Twitter Stream. Note that the UserID refers to ID of a user in Twitter while VenueID refers to ID of a venue in Foursquare. Indeed when such a tuple occurs, it means the user has linked his/her Foursquare account to Twitter account. We have aggregated all the tuples with the same UserID as the historical check-in records of the particular user. We have enriched the dataset by crawling the venue profile information for the venues occurring in the dataset via the public Foursquare Application Programming Interface (API) with VenueID. We have furthermore obtained friendship information among the users via the public Twitter API with UserID.

From the overall collection, we have selected three datasets, which we will refer to as the New York City (NYC) dataset, the California (CA) dataset, and the United States of America (USA) dataset. These datasets were chosen because they allow us to evaluate our method at different geographic scales. Note that we have excluded the New York and California data from the USA dataset, to ensure that all three datasets are disjoint. The check-in records in NYC dataset are located in NYC, which can be regarded as a city-wise dataset; Check-in records in the CA dataset are distributed across several major cities in CA, and thus we can regard this as a state-wise dataset. The USA dataset contains check-in records across the whole country, and thus this can be regarded as a country-wise dataset.

The same processing strategy was adopted for all three datasets. For each user, we first allocated 60% of his/her check-in records to the training set, 20% to the validation set, and 20% to the testing set. The splits between these datasets are chronological. After splitting each dataset, check-in records were enriched by adding companion information and venue profiles. For a particular check-in record associated with a venue, we define companions as friends of the user who have also checked-in at this venue. Although they may not have visited the venue together with the user, they are treated

[1] http://infolab.tamu.edu/data/

Dataset	USA			CA			NYC		
Metric	Acc@1	Pre@5	Rec@5	Acc@1	Pre@5	Rec@5	Acc@1	Pre@5	Rec@5
Core	0.532	0.220	0.719	0.567	0.241	0.732	0.483	0.212	0.689
Core-NoG	0.493	0.184	0.698	0.516	0.213	0.709	0.431	0.172	0.603
VP	0.454	0.183	0.647	0.462	0.207	0.668	0.312	0.154	0.538
RSM	0.382	0.161	0.593	0.426	0.191	0.626	0.249	0.121	0.432
LDA	0.364	0.150	0.607	0.403	0.179	0.655	0.270	0.139	0.471

Table 3: Performance of companion recommendation

as companions so that they can be potentially recommended in the future. Note that companions are generated separately within each split. Moreover, we have extracted category labels of all venues and added them to the check-in records associated with the venues. We then removed those check-in records that have no companions. We also removed users who have less than five check-in records. The statistics of the datasets, after these preprocessing steps is shown in Table 2.

Following previous works on recommendation [45, 27, 17], we evaluate the recommendation performance of our framework using three metrics: Pre@5(Precision at five), Rec@5 (Recall at five), and the average accuracy of the one-companion recommendation task.

The parameters of our model include the number of latent friendship types and the number of latent topics. We make use of the validation dataset to select suitable values for these parameters. In this way, for the NYC, CA and USA datasets, the number of friendship types was selected as 70, 100, and 100, respectively, and the number of latent topics was selected as 100, 150, and 150.

Our full framework is denoted as *CORE* (for Companion Recommender). We also investigate one variant of our framework that does not consider geographical information, denoted as *CORE-NoG*. Specifically, the value of $G(c, v)$ is set to zero for all companions and venues in this variant.

4.2 Comparative Methods

Since companion recommendation is a new task, there are no existing models that directly solve it. Therefore, we have instead adapted some related models to tackle the task of performance comparison, as described below. Specifically, we employ three comparative methods. The first comparative method is treated as a representative of friendship-strength based methods. By comparing with this method, we can investigate the importance of considering the target venue in handling our proposed task. The second method is based on Latent Dirichlet Allocation (LDA). A major difference between our proposed model and LDA is that we model friendship relations between users. This comparison will enable us to analyse the effectiveness of using friendship types. The third method is an intuitive approach that is based on user's preference on venue. Friends that have high preference score on the target venue is recommended as companions. These three comparative methods are described next.

Relationship-Strength Method (RSM): Xiang et al. [43] proposed a model to predict the strength of relationship between two friends, which can be adopted to solve the companion recommendation task. This comparative model infers the friendship strength between two users on LinkedIn and Facebook. For each user pair, it considers two kinds of features: the similarity between the two users and a binary interaction vector. Each entry of the latter vector indicates whether the corresponding interaction exists between the user pair. The interactions include one user views the profile of the other user, recommends the other user, and so on. We adapt their algorithm with features derived from check-in records. In particular, we calculate the similarity value by counting how many times they have checked in at the same venue. Furthermore, interactions are captured by checking whether they ever checked-in at venues with the same category, where each entry in the interaction vector is associated with a particular category. Companion recommendations for each user can be obtained by ranking the user's friends according to their friendship strength. Note that in contrast with our model, this model will generate the same set of recommended companions for each venue.

LDA-Based Method (LDA). This comparison method makes use of LDA to learn the topic relations between users, companions and venues. Given the user u and the venue v, candidate companions c are ranked according to the following score: $\sum_z \theta_{u,z} * \theta_{c,z} * p_v(z)$, where $\theta_{u,z}$ and $\theta_{c,z}$ are the posterior distributions over topics z for u and c, representing the interest of the user u and the user c respectively, and $p_v(z)$ is the posterior distribution of venues for the topic z, representing the topics of the venue.

Venue-Preference Method (VP). A score evaluating the preference is calculated for each pair of user and venue. To tackle the companion recommendation problem, where a query user and a target venue are given, companions are ranked by the preference score to the target venue. We adapt a state-of-the-art POI recommendation model [17] to compute the user's preference on venues.

4.3 Quantitative Results

The performance of our framework and the comparative methods on the testing dataset is shown in Table 3. The results show that the performance of our framework is consistently and substantially better than the comparative methods across all the datasets. The differences between our model, on the one hand, and the Relationship-Strength method, LDA-based method, and Venue-Preference method, on the other hand, are statistically significant for all datasets and all evaluation metrics, based on the paired t-test with $p < 0.01$.

The Relationship-Strength method performs substantially worse than our framework. The main underlying reason is that it recommends companions based on the strength of the relationship between two users, regardless of the given input venue. Although this method can infer stronger ties between users who have more check-ins at the same venue and more similarity in profiles, it cannot exploit information about the target venue when making companion recommendations. We observe from the results that the incorporation of geographical information would generally improves the performance, which is consistent with what is found in most LBSN-based tasks. When removing the geographical information from our framework, it still performs better than the LDA-based method. The major difference between these two methods is again that our framework makes companion recommendations with suitable friendship types that match the characteristics of the target venue. This confirms our hypothesis that incorporating friendship types results in better companion recommendations. Our frame-

Friendship Type	Category Labels
1	College Stadium, Convenience Store, College Hockey Rink, Rock Climbing Spot, Climbing Gym, Yoga Studio, College Basketball Court, College Gym, College Auditorium, Synagogue
2	Art Gallery, Art Museum, Museum, Park, Music Venue, Hotel, History Museum, Non-Profit, Concert Hall, Opera House
3	Asian Restaurant, Japanese Restaurant, Sushi Restaurant, Ramen/Noodle House, Thai Restaurant, Chinese Restaurant, Korean Restaurant, Sandwich Place, American Restaurant, Vietnamese Restaurant
4	Gym, Gym/Fitness Center, Athletics & Sports, Event Space, Community Center, Yoga Studio, Pool, College Gym, College Rec Center, Tanning Salon
5	Office, Coworking Space, Event Space, Tech Startup, Building, Advertising Agency, Professional & Other Places, Conference Room, General Entertainment, Design Studio
6	Event Space, Conference Room, Office, Wine Bar, Winery, General Entertainment, Convention Center, Performing Arts Venue, Vineyard, Cafeteria

Table 4: Semantic representation of friendship types

Friendship Type	Names of Venues Visited
1	Hoover Tower, STAPLES Center, Rose Bowl Stadium, Clancey's Market & Deli, 7-Eleven
2	Last Rites Tattoo Theatre and Art Gallery, Japan Society, Art Directors Club
3	Mingalaba Restaurant, Carnitas' Snack Shack, CUCINA urbana, Sol Food Puerto Rican Cuisine
4	Fit Athletic Club, Chuze Fitness, Equinox, UCSF Bakar Fitness & Rec Center
5	Tesla Motors HQ, Foursquare SF, Festival Pavilion, Yahoo!
6	Yahoo!, Zero Zero, The Strand

Table 5: Venues that pairs of users with corresponding friendship types both visited

work is also better than the method based on venue preference, which is adapted from a POI recommendation system.

4.4 Qualitative Analysis

4.4.1 Friendship Type Analysis

Because the category labels of venues are taken into account when modeling friendship types, we can use our framework to produce a semantic description of the latent friendship types. In particular, the parameter Φ^l inferred from our model encodes the relevance of friendship types to categories. To illustrate this, we present some friendship types and their most relevant categories in Table 4. From the category labels describing the corresponding friendship types, we can clearly understand the nature of these friendship types. Intuitively, the friendship types 1 to 5 in Table 4 correspond, respectively, to schoolmates, friends interested in art, friends interested in Asian food, friends interested in sports, co-workers. Friendship type 6, which is slightly different with friendship types 1 to 5, corresponds to co-workers who are also interested in entertainment after work.

Our framework can also characterize the latent friendship types via the inferred parameter Φ^c. To illustrate this, we have selected some user pairs whose friendship types match the ones listed in Table 4. For these user pairs, we then analyze which venues they have most frequently visited. Several of these top venues are listed in Table 5, where the friendship types match those from 4 line by line. We can observe from Table 5 that the venues which the user pairs have ever visited are indeed tightly related to the friendship types described in Table 4. This further illustrates the effectiveness of our framework in characterizing the friendship types between users.

4.5 Sensitivity Analysis

4.5.1 Effect of Parameters

We evaluate the performance of our framework under different settings of the number of friendship types and the number of topics as depicted in Figure 3. When we vary the number of friendship types, we keep the number of topics fixed at 100. Similarly, when we vary the number of topics, we keep the number friendship types fixed at 20. The results show that the performance of our framework is rather robust to changes in the number of friendship types and the number of topics, for the three considered datasets. The performance generally increases as the corresponding parameters increase and becomes stable after a certain point.

4.5.2 Convergence Analysis

We employ the Gibbs EM method for the inference of parameters. The convergence behavior is shown in Figure 5. It shows that the performance of our framework generally gets stable after 700 iterations for the NYC and CA datasets. For USA dataset, it takes about 1000 iterations for the performance to converge.

5. CONCLUSION

We have proposed a framework to solve the problem of recommending companions to users, given a particular venue that the user is interested in visiting. Companion recommendations are made by learning the relationship between latent venue topics and latent friendship types, between latent venue topics and the previous check-in behaviour of users, and between latent friendship types and the categories of previously visited venues. Experimental results show that our framework can solve this task effectively.

6. ACKNOWLEDGMENTS

The work described in this paper is supported by a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14203414). This work is also supported by ERC Starting Grant 637277. Another support is from Microsoft Research Asia Urban Informatics Grant FY14-RES-Sponsor-057.

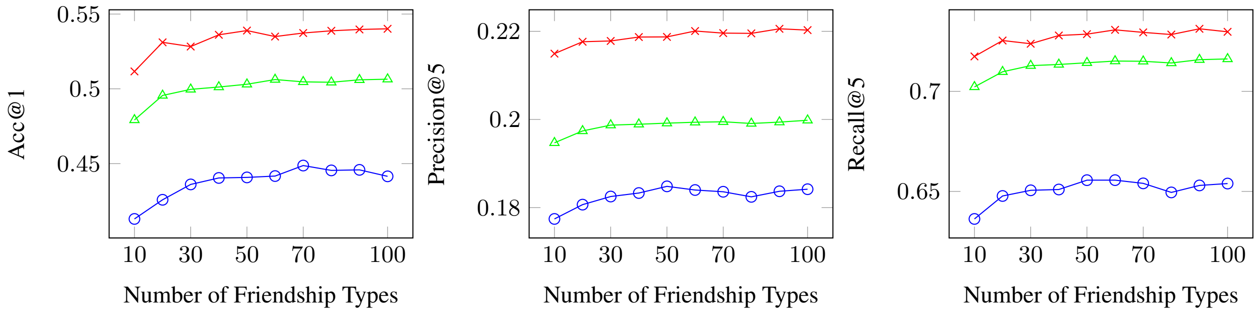

Figure 3: Effect of varying the number of friendship types in our model. —△— denotes results for USA, —×— denotes for CA, and —○— denotes for NYC.

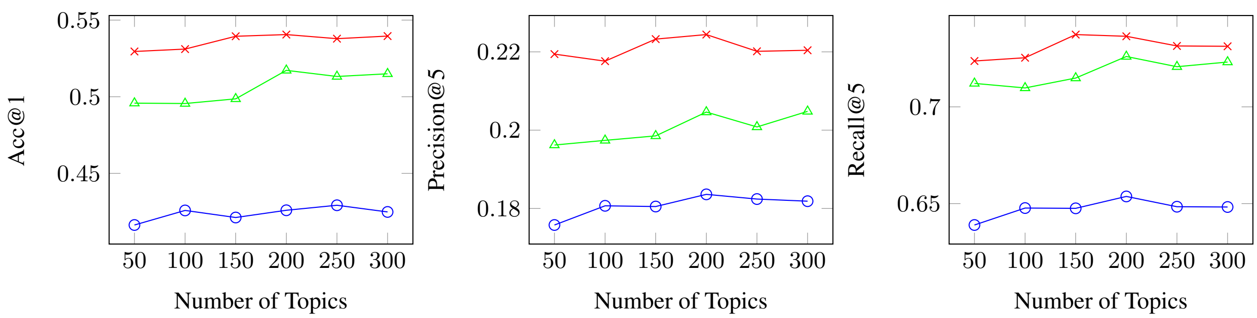

Figure 4: Effect of varying the number of topics in our model. —△— denotes results for USA, —×— denotes for CA, and —○— denotes for NYC.

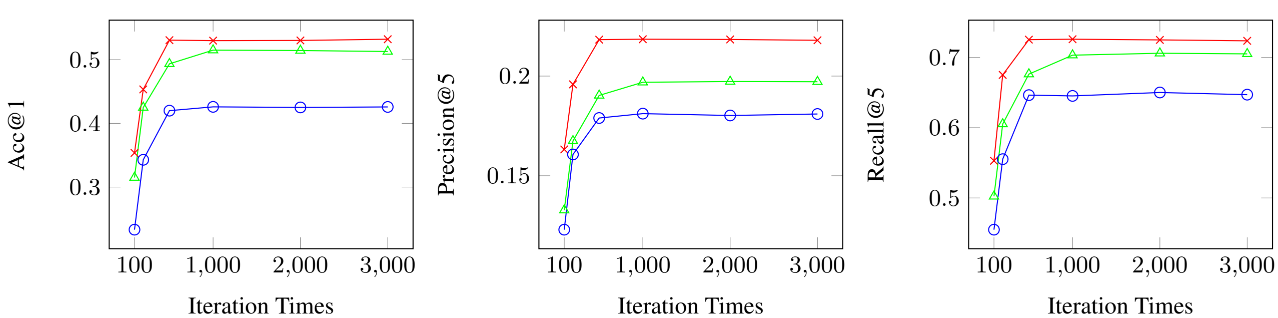

Figure 5: Effect of iteration times in our model. —△— denotes results for USA, —×— denotes for CA, and —○— denotes for NYC.

7. REFERENCES

[1] L. Backstrom and J. Leskovec. Supervised random walks: predicting and recommending links in social networks. In *WSDM*, pages 635–644, 2011.

[2] Z. Burda, J. Duda, J. Luck, and B. Waclaw. Localization of the maximal entropy random walk. *Physical Review Letters*, 102(16):160602, 2009.

[3] J. Cheng, T. Yuan, J. Wang, and H. Lu. Group latent factor model for recommendation with multiple user behaviors. In *SIGIR*, pages 995–998, 2014.

[4] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui. Exploring millions of footprints in location sharing services. *ICWSM*, 2011:81–88, 2011.

[5] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: User movement in location-based social networks. In *KDD*, pages 1082–1090, 2011.

[6] P. De Meo, E. Ferrara, G. Fiumara, and A. Provetti. On facebook, most ties are weak. *Communications of the ACM*, 57(11):78–84, 2014.

[7] J. Eisenstein, A. Ahmed, and E. P. Xing. Sparse additive generative models of text. In *ICML*, 2011.

[8] S.-S. Feng, X.-T. Li, Y.-F. Zeng, C. Gao, Y.-M. Chee, and Q. Yuan. Personalized ranking metric embedding for next new poi recommendation. In *IJCAI*, 2015.

[9] H. Gao, J. Tang, X. Hu, and H. Liu. Exploring temporal effects for location recommendation on location-based social networks. In *RecSys*, pages 93–100, 2013.

[10] H. Gao, J. Tang, X. Hu, and H. Liu. Content-aware point of interest recommendation on location-based social networks. AAAI, 2015.

[11] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *SIGCHI*, pages 211–220. ACM, 2009.

[12] A. Gionis, T. Lappas, K. Pelechrinis, and E. Terzi. Customized tour recommendations in urban areas. In *WSDM*, pages 313–322, 2014.

[13] J. Hannon, M. Bennett, and B. Smyth. Recommending twitter users to follow using content and collaborative filtering approaches. In *RecSys*, pages 199–206, 2010.

[14] G. Heinrich. Parameter estimation for text analysis. *University of Leipzig, Tech. Rep*, 2008.

[15] B. Hu and M. Ester. Spatial topic modeling in online social media for location recommendation. In *RecSys*, pages 25–32, 2013.

[16] I. Kahanda and J. Neville. Using transactional information to predict link strength in online social networks. In *ICWSM*, pages 74–81, 2009.

[17] T. Kurashima, T. Iwata, T. Hoshide, N. Takaya, and K. Fujimura. Geo topic model: Joint modeling of user's activity area and interests for location recommendation. In *WSDM*, pages 375–384, 2013.

[18] K. W.-T. Leung, D. L. Lee, and W.-C. Lee. Clr: a collaborative location recommendation framework based on co-clustering. In *SIGIR*, pages 305–314, 2011.

[19] H. Li, R. Hong, S. Zhu, and Y. Ge. Point-of-interest recommender systems: A separate-space perspective. In *ICDM*, pages 231–240, 2015.

[20] M. Li, A. Ahmed, and A. J. Smola. Inferring movement trajectories from GPS snippets. In *WSDM*, pages 325–334, 2015.

[21] R.-H. Li, J. X. Yu, and J. Liu. Link prediction: the power of maximal entropy random walk. In *CIKM*, pages 1147–1156, 2011.

[22] X. Li, G. Cong, X.-L. Li, T.-A. N. Pham, and S. Krishnaswamy. Rank-geofm: A ranking based geographical factorization method for point of interest recommendation. In *SIGIR*, pages 433–442, 2015.

[23] D. Lian, Y. Ge, F. Zhang, N. J. Yuan, X. Xie, T. Zhou, and Y. Rui. Content-aware collaborative filtering for location recommendation based on human mobility data. In *ICDM*, pages 261–270, 2015.

[24] D. Lian, C. Zhao, X. Xie, G. Sun, E. Chen, and Y. Rui. GeoMF: Joint geographical modeling and matrix factorization for point-of-interest recommendation. In *KDD*, pages 831–840, 2014.

[25] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *CIKM*, pages 556–559, 2003.

[26] K. H. Lim, J. Chan, C. Leckie, and S. Karunasekera. Personalized tour recommendation based on user interests and points of interest visit durations. In *IJCAI*, 2015.

[27] B. Liu, Y. Fu, Z. Yao, and H. Xiong. Learning geographical preferences for point-of-interest recommendation. In *KDD*, pages 1043–1051, 2013.

[28] D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.

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

[30] T. Lou, J. Tang, J. Hopcroft, Z. Fang, and X. Ding. Learning to predict reciprocity and triadic closure in social networks. *TKDD*, 7(2):5, 2013.

[31] Z. Qiao, P. Zhang, Y. Cao, C. Zhou, L. Guo, and B. Fang. Combining heterogenous social and geographical information for event recommendation. In *AAAI*, 2014.

[32] D. M. Romero and J. M. Kleinberg. The directed closure process in hybrid social-information networks, with an analysis of link formation on twitter. In *ICWSM*, 2010.

[33] I. Ronen, I. Guy, E. Kravi, and M. Barnea. Recommending social media content to community owners. In *SIGIR*, pages 243–252, 2014.

[34] A. Sadilek, H. Kautz, and J. P. Bigham. Finding your friends and following them to where you are. In *WSDM*, pages 723–732, 2012.

[35] A. Salehi-Abari and C. Boutilier. Preference-oriented social networks: Group recommendation and inference. In *RecSys*, pages 35–42, 2015.

[36] R. Schifanella, A. Barrat, C. Cattuto, B. Markines, and F. Menczer. Folks in folksonomies: social link prediction from shared metadata. In *WSDM*, pages 271–280, 2010.

[37] S. Sintos and P. Tsaparas. Using strong triadic closure to characterize ties in social networks. In *KDD*, pages 1466–1475, 2014.

[38] D. Song, D. A. Meyer, and D. Tao. Top-k link recommendation in social networks. In *ICDM*, pages 389–398, 2015.

[39] K. Subbian, A. Banerjee, and S. Basu. PLUMS: predicting links using multiple sources. In *SDM*, pages 370–378, 2015.

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

[41] S. Wan, Y. Lan, J. Guo, C. Fan, and X. Cheng. Informational friend recommendation in social media. In *SIGIR*, pages 1045–1048, 2013.

[42] W. Wang, H. Yin, L. Chen, Y. Sun, S. Sadiq, and X. Zhou. Geo-sage: A geographical sparse additive generative model for spatial item recommendation. In *KDD*, pages 1255–1264.

[43] R. Xiang, J. Neville, and M. Rogati. Modeling relationship strength in online social networks. In *WWW*, pages 981–990, 2010.

[44] G. Yuan, P. K. Murukannaiah, Z. Zhang, and M. P. Singh. Exploiting sentiment homophily for link prediction. In *RecSys*, pages 17–24, 2014.

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

[46] Q. Yuan, G. Cong, Z. Ma, A. Sun, and N. M. Thalmann. Time-aware point-of-interest recommendation. In *SIGIR*, pages 363–372, 2013.

[47] J. Zhang, X. Kong, and P. S. Yu. Transferring heterogeneous links across location-based social networks. In *WSDM*, pages 303–312, 2014.

[48] J. Zhang, C. Wang, P. S. Yu, and J. Wang. Learning latent friendship propagation networks with interest awareness for link prediction. In *SIGIR*, pages 63–72, 2013.

[49] J. Zhang, P. S. Yu, and Z.-H. Zhou. Meta-path based multi-network collective link prediction. In *KDD*, pages 1286–1295, 2014.

[50] J.-D. Zhang and C.-Y. Chow. Geosoca: Exploiting geographical, social and categorical correlations for point-of-interest recommendations. In *SIGIR*, pages 443–452, 2015.

[51] W. Zhang and J. Wang. A collective bayesian poisson factorization model for cold-start local event recommendation. In *KDD*, pages 1455–1464, 2015.

[52] W. Zhang, J. Wang, and W. Feng. Combining latent factor model with location features for event-based group recommendation. In *KDD*, pages 910–918, 2013.

[53] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang. Collaborative location and activity recommendations with GPS history data. In *WWW*, pages 1029–1038, 2010.

[54] Y. Zhong, N. J. Yuan, W. Zhong, F. Zhang, and X. Xie. You are where you go: Inferring demographic attributes from location check-ins. In *WSDM*, pages 295–304, 2015.

[55] W.-Y. Zhu, W.-C. Peng, L.-J. Chen, K. Zheng, and X. Zhou. Modeling user mobility for location promotion in location-based social networks. In *KDD*, pages 1573–1582, 2015.