PGT: Measuring Mobility Relationship using Personal, Global and Temporal Factors

Hongjian Wang*, Zhenhui Li*, Wang-Chien Lee†
*College of Information Sciences and Technology, Pennsylvania State University
†Department of Computer Science and Engineering, Pennsylvania State University
*{hxw186, jessieli}@ist.psu.edu, †wlee@cse.psu.edu

***Abstract*—Rich location data of mobile users collected from smartphones and location-based social networking services enable us to measure the mobility relationship strength based on their interactions in the physical world. A commonly-used measure for such relationship is the frequency of meeting events (i.e., co-locate at the same time). That is, the more frequently two persons meet, the stronger their mobility relationship is. However, we argue that *not all the meeting events are equally important in measuring the mobility relationship* and propose to consider *personal* and *global* factors to differentiate meeting events. Personal factor models the probability for an individual user to visit a certain location; whereas the global factor models the popularity of a location based on the behavior of general public. In addition, we introduce the *temporal factor* to further consider the time gaps between meeting events. Accordingly, we propose a unified framework, called PGT, that considers personal, global, and temporal factors to measure the strength of the relationship between two given mobile users. Extensive experiments on real datasets validate our ideas and show that our method significantly outperforms the state-of-the-art methods.**

***Keywords*-mobility; relationship strength; spatiotemporal; social computing**

I. Introduction

Rapid advances in positioning technology and popularity of mobile devices facilitate easy collection of rich location information. Also thanks to Web 2.0 technology, many social media are now available for their users to share opinions, pictures and even footprints. For example, mobile users can post geo-tagged microblogs on Twitter, photos with location information on Flickr, and check-ins on Foursquare. The ever-growing location data provides us with an opportunity to study the human behavior in the *physical* world.

One interesting question is whether we can measure the *mobility relationship* between two mobile users based on their interactions in the physical world (e.g., how often two persons meet and where and when they meet). In contrast to the *cyber* (i.e., online) relationships, the mobility relationships are measured based on the movement data. Understanding such a mobility relationship has a profound impact in advancing social sciences and a great benefit to a number of applications such as transportation scheduling [1], urban planning [2], recommendation [3], [4], advertisement targeting [5], privacy protection [6], and anomaly detection [7].

In this paper, we study the problem of measuring the mobility relationship strength between two users based on their location history data.[1] Conventionally, the mobility relationship has been measured by *meeting frequency*, i.e., how often two persons co-locate at the same time. However, we argue that *these meeting events should not be treated equally*, and propose to consider the following factors in weighting the importance of meeting events.

Personal factor. The same location carries different meanings for different persons. For example, the Times Square in New York could be a travel destination for a tourist but also an office neighborhood for a person working there. For two persons whose offices are located near Times Square, the number of co-locating events between them could be high, even though such meeting events do not necessarily indicate that they know each other. On the other hand, for two persons who only visit Times Square three times but always co-locate there at the same time, it is likely that they have a strong relationship and therefore meet each other there with some intentions, e.g. for business or get-together. In summary, the importance of a meeting event needs to *consider an individual user's probability to visit a certain location*, which is termed as the *personal factor* in this paper.

Global factor. Different from the personal factor which considers *an individual's probability* to visit a location, the *global factor* captures the popularity of a location to the general public (which can be estimated using *all the users* in the dataset). Some locations are frequently visited by many people, such as the downtown in the city and a big football stadium, whereas other locations are more specific to only a few people, such as a private house. In a popular public place, it is more likely for two strangers to co-locate by coincidence. Thus, such meeting events are less indicative for a relationship. In contrast, a meeting event at a private place often indicates a stronger mobility relationship.

Temporal factor. Temporal factor considers the time gaps between consecutive meeting events. For example, suppose two users attended the same football game or traveled on the same train and they both made several check-ins during that period. Then, they are likely to have several co-locating events although they may not know each other. On the other hand, meeting events which are separated by a long time period

[1] In this paper, we measure the users' relationships solely based on their physical interactions. It is also possible to mine the relationships from other types of interactions (e.g., phone calls and online communications), but it is beyond the scope of this paper.

(e.g., 10 days) are less likely to be coincident. Therefore, such events should have more weights in determining the mobility relationship strength than events which occurred within a short time interval (e.g., one hour).

In the literature, meeting frequency has been widely used to measure the strength of mobility relationship between moving objects [8], [9], [10], [11], [12]. Recent works [13], [14] have considered the global factor of a location by using the entropy to measure its popularity. A meeting event is then weighted based on the entropy of the meeting location (i.e., visited by many different users). However, *few* works consider the personal factor and the temporal factor, which are also important in differentiating meeting events. For example, as Times Square is a popular location, all the meeting events there are likely to be treated as coincidences by the global factor. However, from the viewpoint of the proposed personal factor, the "popularity" of a location should be interpreted differently for a New Yorker and a tourist, because a tourist has a much lower probability visiting there than a New Yorker. In addition, two consecutive meeting events which occur at Times Square in the same hour often carry different meanings than two events which occur in two different weekends. Such differences can be captured by the proposed personal and temporal factors.

In this paper, we propose a unified framework, called `PGT`, that incorporates all the aforementioned factors to measure the mobility relationship. The framework consists of two components. (i) Background modeling – here we extract the personal background for each user and the global background for a location from a dataset of mobile users' location history. A density-based method is used to model the personal background and an entropy-based method is used to model the global background. (ii) Mobility relationship mining – by incorporating both the personal factor of each user and the global factor of each location, we determine the weights of the meeting events and the strength of mobility relationship. In addition, we further incorporate the temporal factor to penalize the consecutive meetings in a short time interval and to reward the sporadic meetings with longer time span.

The major contributions of this paper are as follows. **1.** We study two new and important factors, the personal factor and the temporal factor, in mobility relationship mining. To the best of our knowledge, these two factors have not been examined in the literature. **2.** We propose a unified framework to measure the strength of mobility relationships, which combine the personal, global, and temporal factors. We show that these factors complement each other under our framework. **3.** We conduct extensive experiments on real data to demonstrate the effectiveness of our method. We provide in-depth analyses of the experiment results and give great insights into the mobility relationship mining problem.

The remaining of the paper is organized as follows. Section II formally states our relationship mining problem and Section III outlines the general framework of our approach. Section IV and Section V describe each step in our framework in detail. Section VI reports our experimental results. Section VII discusses related work. Finally, Section VIII concludes our study.

II. Problem Definition

Given a location dataset of n users, the location history of a user i is represented as a sequence of locations and timestamps: $S_i = \langle (loc_1^i, t_1^i), (loc_2^i, t_2^i), \ldots, (loc_m^i, t_m^i) \rangle$, where each location is a geographic coordinate. Given any two users i and j, our goal is to calculate a *relationship measure* F_{ij} based on their mobility data, i.e. meeting events.

A **meeting event** is formed when two users being spatially close at the same time. Specifically, for a pair of location records $(loc_p^i, t_p^i) \in S_i$ and $(loc_q^j, t_q^j) \in S_j$, if the location distance is less than a distance threshold δ and the time difference is less than a time threshold τ (i.e., $dist(loc_p^i, loc_q^j) < \delta$ and $|t_p^i - t_q^j| < \tau$), this pair of records forms a meeting event. Here, δ and τ are two application-dependent parameters and can be set experimentally. In this paper, we fix $\delta = 30(meters)$ and $\tau = 1(hour)$.

Let $E_{ij} = \{e_1, e_2, \cdots\}$ denote the sequence of meeting events between users i and j. Each event e_k contains a location and a timestamp, $e_k = (loc_k, t_k)$, where the values are simply computed by averaging the locations and timestamps of the corresponding pair of location records of the two users.

The relationship measure F_{ij} is a function of all the meeting events between users i and j:

$$F_{ij} = G(E_{ij}). \tag{1}$$

Our goal is therefore to find a discriminative function G which well differentiates the mobility relationship strength. In practice, a simple implementation of G could be the frequency of meeting events (i.e., **meeting frequency**):

$$G_0(E_{ij}) = |E_{ij}|. \tag{2}$$

However, as we discussed before, the meeting frequency treats meeting events equally, so it is not sufficient as a measure of the relationship strengths.

III. Framework

The *mobility relationship strength* between a pair of users is decided by the set of meeting events between them. We argue that different meeting events contribute differently in determining the relationship between users. In this paper, we propose a unified framework to explore three factors (personal, global, and temporal factors), as shown in Figure 1. The framework consists of two major components: *background modeling* and *relationship mining*. We first give a high-level overview of the components here and then discuss the technical details in Section IV and Section V, separately.

Background Modeling (details in Section IV). In the first component, we model the backgrounds for each individual user and for the locations. In particular, as shown in the upper part of Figure 1, there are two backgrounds: the personal background and the global background. The background modeling is independent from Component 2 on relationship mining as

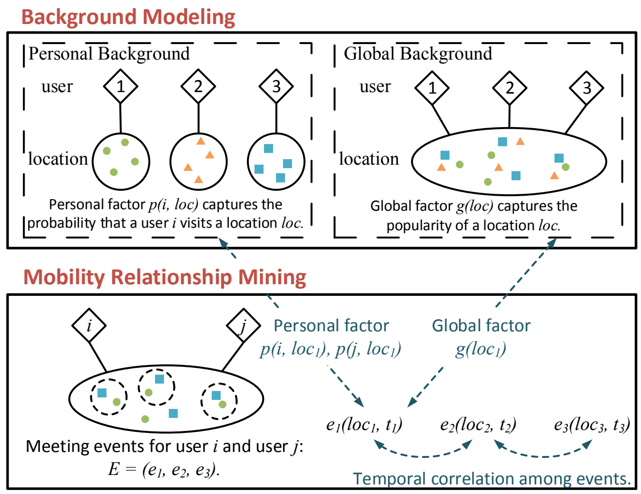

Fig. 1. The PGT framework for mobility relationship mining.

the background modeling does not look at the meeting events between a pair of users.

Personal background (Section IV-A). A user often goes to some places (e.g., his office) much more often than the other places (e.g., a bookstore), and different locations often carry different meanings for a specific user. Therefore, it is important to look into this user's location history and model the importance of each place to the user. In this paper, we call such a model the *personal mobility background* of the user. To model the personal mobility background, we propose to estimate the probability $\rho(i, loc_k)$ of a user i visiting any location loc_k on the map.

Global background (Section IV-B). In addition to modeling the personal usage of a location for each individual user, we also notice that, if we aggregate the location history of all the users, different locations may exhibit different characteristics, such as *popularity*. For example, a shopping center is much more popular than someone's home, and thus it is visited by a much larger number of people. Therefore, it is important to treat these places differently when inferring the mobility relationship strength of users. To model the global mobility background, we capture the popularity $g(loc_k)$ of a location loc_k using the entropy based on the number of visits by each user to this location.

Mobility Relationship Mining (details in Section V). In the second component, we focus on mining the mobility relationship between *a pair of users* i *and* j, using the set of their meeting events, which are weighted based on the personal factor and the global factor extracted in the first component. In addition, we also consider the temporal correlation between the meeting events in our mobility relationship strength measure.

Personal factor (Section V-A). The personal background has been modeled as the probability for a user to visit a location. On the other hand, if a meeting event occurs at some place where one or both users frequently visit, this event is more likely to be a *coincidence* and therefore should carry less weight. If two users i and j meet at a location they both rarely visit (e.g., a travel destination or a restaurant far away from his home and office), they are more likely to be true friends. Therefore, such meeting event serves as a strong indicator of their relationship and should be rewarded.

Global factor (Section V-B). The global background has modeled the popularity of a location. If two users meet at a popular location (e.g., a shopping center) where many people co-locate, such a meeting event should carry less weight. On the other hand, a meeting event at a less popular location (e.g., a private house) should carry more weight.

Temporal factor (Section V-C). In addition to computing the personal and global factors for each meeting event, we further consider the time gaps between two events. The idea is to penalize the meeting events with very short time gaps (e.g., meeting 10 times in a day) and to reward the ones which span a long period (e.g., meeting once every month for 10 months).

IV. Background Modeling

Following the overall framework introduced earlier, in this section, we detail how to capture the user's personal preference to each location (i.e., the personal background) and the popularity of a location over the entire population (i.e., the global background).

A. Modeling Personal Mobility Background

The personal mobility background models the probability that a user visits a location. The probability for user i to visit a location loc can be computed by:

$$p(i, loc) = \frac{|\{(loc_k, t_k) \in S_i : loc_k \sim loc\}|}{|S_i|}, \tag{3}$$

where $\sim$ is the equivalence relation denoting that two elements are equivalent.

To judge the equivalence relation of two locations, two methods are frequently used. One is to partition the space into grids and then judge the equivalence based on their corresponding grid IDs. Another way is to judge whether two locations are within a certain threshold. But both methods pose hard constraints on the equivalence test. Suppose two geographical records are close to each other, but they just happen to be partitioned into two different grids or their distance is slightly larger than the threshold, then they will not be properly considered in the probability calculation.

To overcome the above issue, we propose a density function to model the personal mobility background. The proposed model naturally incorporates the distances from any point to all the recorded locations without using any hard constraint. Specifically, for a location loc_k on the map, we model its density with respect to user i by considering the distances between loc_k and all the recorded locations of user i:

$$\rho(i, loc_k) = \sum_{(loc_q^i, t_q^i) \in S_i} \exp(-c_d \cdot dist(loc_k, loc_q^i))/|S_i|, \tag{4}$$

where c_d is a parameter and $\rho(i, loc_k)$ is the density of location. Although $\rho(i, loc_k)$ is not a probability value but a density value, it captures the same semantics as the user's probability to visit a place. For example, if a user visits loc_k

less frequently, then loc_k must be far away from most of the locations visited by user i and thus has a low density.

In our density model, the parameter c_d determines how fast the impact (or correlation) of a recorded location on its neighborhood falls as the distance increases. In practice, the impact of one location on another location is closely related to human's actual mobility. Intuitively, two locations should have a big impact on each other if they are within walking distances ($< 1km$), and have a near-zero impact if driving is required ($> 5km$). We further examine the sensitivity of our method to the parameter c_d in Section VI-C.

B. Modeling Global Mobility Background

The global background captures the *location popularity* inferred from all the mobile users. If a location is frequently visited by a large number of mobile users, it could be a public place, such as a tourist spot or a train station. Meeting events at such locations are often less indicative of the strength of mobility relationship. On the other hand, if two people meet at a less popular location, such as a private house, they are more likely to have a real relationship.

To model the popularity of a place, the *location entropy* is proposed in [15] and also adopted in recent work [14]. Specifically, let $S_i(loc_k)$ be the set of location records of user i visiting location loc_k,

$$S_i(loc_k) = \{(loc_q^i, t_q^i) \in S_i : loc_q^i \sim loc_k\}, \tag{5}$$

where condition $loc_q^i = loc_k$ can be judged either by the location IDs or by the distance between two locations. Then, the probability for user i to visit location loc_k is

$$P(i, loc_k) = |S_i(loc_k)| / \sum_i |S_i(loc_k)|. \tag{6}$$

Note that the probability $P(i, loc_k)$ is different from the density function $\rho(i, loc_k)$ in our personal background model. Here, $P(i, loc)$ is the ratio of user i's visits to location loc over all the visits from the entire population, whereas $\rho(i, loc)$ captures the ratio of one user's visits to a location over the size of his own location history.

The Shannon entropy of a location loc_k can be estimated using the probability vector of all users visiting this location:

$$g(loc_k) = - \sum_{i:P(i,loc_k)\neq 0} P(i, loc_k) \cdot \log P(i, loc_k). \tag{7}$$

Here, a low entropy implies that a location is visited by few users and a high entropy indicates that a location is visited by many different users. For example, if a location is visited by only one user, the entropy of this location is $g(loc_k) = -1 \log 1 = 0$. Meanwhile, if the location is visited by all the users with an equal probability, we then have $g(loc_k) = -\sum_{i=1}^{n} \frac{1}{n} \log \frac{1}{n} = \log(n)$.

V. Mining Mobility Relationship

In this section, we focus on measuring the relationship strength between two users i and j. As discussed in Section II, the relationship strength is a function $G(E_{ij})$ of the set of meeting events $E_{ij} = \{e_1, e_2, \cdots\}$. We will consider both personal factor and global factor in weighting the meeting events. In addition, we will incorporate the temporal correlations between the events into function $G(E_{ij})$.

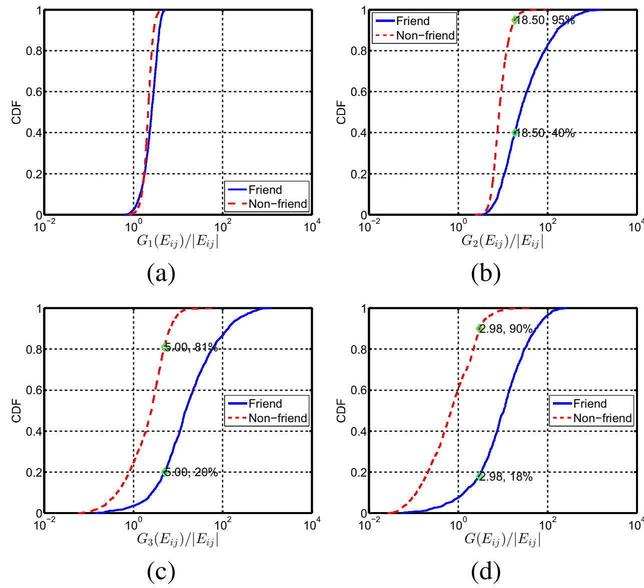

Fig. 2. Comparison of the cumulative distribution functions of various weights for friend and non-friend pairs. (a) Average personal weight. (b) Maximum personal weight. (c) Combination of personal and global weights. (d) Combination of personal, global, and temporal correlation weights. The bigger the area between the two curves is, the more effective the measure is in differentiating friends and non-friends.

A. Personal Factor

Given the density model of each user, we aim to determine the significance of a meeting event e_k between users i and j at location loc_k. As we mentioned before, if they meet at a place where they frequently visit, then this meeting event is likely to happen by chance; if two people meet at a place where they rarely visit, then they are likely to have a real relationship. Therefore, for a meeting event $e_k \in E_{ij}$, we define the personal factor weight of this meeting event as follows:

$$w_{ij}^p(e_k) = -\log(\rho(i, loc_k) \cdot \rho(j, loc_k)). \tag{8}$$

Accordingly, for the set of meeting events E_{ij} between users i and j, we can calculate their relationship strength as follows:

$$G_1(E_{ij}) = \sum_{e_k \in E_{ij}} w_{ij}^p(e_k) = \overline{w}_{ij}^p \times |E_{ij}|, \tag{9}$$

where $\overline{w}_{ij}^p$ is the average weight across all meeting events.

To verify the effectiveness of $G_1(E_{ij})$ in differentiating mobility relationship, we plot the cumulative distribution function (CDF) of $\overline{w}_{ij}^p$ for all the friend and non-friend pairs in Figure 2(a) using the Gowalla dataset (see Section VI for detailed descriptions of the dataset). Interestingly, we see that two curves almost overlap, which indicates that the average personal weight does not work as well as we expected in separating the friend and non-friend pairs. In other words, it is difficult to use G_1 to differentiate the friend and non-friend pairs in the dataset.

This observation motivates us to further look into the behaviors of the friend pairs, and we realize that many friend pairs do go to places where they rarely visit to meet each other. However, they may also meet at locations where they visit a lot, like a school, an office or a restaurant. By using the average weight in our strength measure, we over penalize the meeting events at rarely visited locations.

Inspired by the above discussion, we propose to explore maximum personal weight over all the meeting events instead, and revise our relationship strength measure accordingly:

$$G_2(E_{ij}) = \max_{e_k \in E_{ij}} \{w^p_{ij}(e_k)\} \times |E_{ij}|. \quad (10)$$

Figure 2(b) shows the CDF curves of the maximum weight for friend and non-friend pairs. Compared with Figure 2(a), we can see that the difference between the two curves in Figure 2(b) becomes much larger. In particular, 60% of friends have weights higher than 18 and only 5% of non-friends have weights higher than the same number. The largest distribution gap in Figure 2(b) is 55%. This illustrates that the maximum personal background weight is indeed an effective measure of friend relationships.

B. Global Factor

For a meeting event $e_k = \{loc_k, t_k\}$, we now consider the factor of the global background. Generally, we wish to reward the meeting events at private locations with low entropies (i.e., small $g(loc_k)$ value) and penalize those at public locations high entropies (i.e., large $g(loc_k)$ value). Following [14], we use the exponential function of the location entropy as the global background weight of an event:

$$w^g_{ij}(e_k) = \exp(-g(loc_k)). \quad (11)$$

One issue with this global factor is that it requires location history data from a large population in order to obtain an accurate estimation of the location popularity. In contrast, the personal factor can be computed from the location history data of two users.

To combine the weights from personal factor and from global factor, we replace the meeting frequency $|E_{ij}|$ in Eq. (10) with the sum of global background weights over all meeting events to obtain a new measure as follows:

$$G_3(E_{ij}) = \max_{e_k} \{w^p_{ij}(e_k)\} \times \sum_{e_k} w^g_{ij}(e_k). \quad (12)$$

Here, we choose the product of the personal and global weights rather than the sum because the weights of these two factors have very different ranges of values. Compared with the sum, the product is much less sensitive to the difference in scale.

In Figure 2(c), we show the CDF curves for the combination of personal and global weights. Compared with Figure 2(b), the difference between the two curves in Figure 2(c) is even bigger. The largest gap increases to 61% from 55%, which indicates that combining these two factors helps us further differentiate the friend and non-friend pairs.

C. Temporal Correlation of Meeting Events

Another important factor in relationship mining is to consider the correlation of the meeting events. Specifically, an event e_k in the set of meeting events E_{ij} between users i and j should be penalized if there are other events that are temporally close to event e_k. The strength of the temporal link between events e_p and e_k can also be modeled by an exponential function with respect to the time difference:

$$l_T(e_k, e_p) = \exp(-c_t \cdot \mid t_k - t_p \mid). \quad (13)$$

Similar to the spatial parameter c_d, the parameter c_t should be chosen based on the correlations of real meeting events. Intuitive, two events that occur in the same hour are often highly correlated, whereas two events that are separated by many days should be considered as independent events.

Next, we define the temporal correlation weight of event e_k based on its temporal link to the previous event:

$$w^t_{ij}(e_k) = \begin{cases} 1 - l_T(e_k, e_{k-1}), & k > 1 \\ 1, & k = 1 \end{cases} \quad (14)$$

Finally, we adjust our relationship strength measure by incorporating the temporal correlation weight as follows:

$$G(E_{ij}) = \max_{e_k} \{w^p_{ij}(e_k)\} \times \sum_{e_k} \Big(w^g_{ij}(e_k) \times w^t_{ij}(e_k) \Big). \quad (15)$$

Figure 2(d) plots the CDF curves of our final relationship measure which combines the personal weight, the global weight and the temporal correlation weight. Compared with Figure 2(c), Figure 2(d) shows an even bigger gap (a maximum distribution gap of 72%) between the two curves, suggesting that the combination can best differentiate the friend and non-friend pairs in practice.

VI. Experiments

In this section, we present a comprehensive performance study of the proposed method on two real datasets. All the experiments are conducted on a 3.4 GHz Intel Core i7 system with 16 GB memory.

A. Dataset and Metrics

We use two different datasets collected from two location-based social networking services, Gowalla and Brightkite [16]. Users share their locations by checking in at places. The check-in records of Gowalla dataset are collected from February 2009 to October 2010, whereas the Brightkite dataset is collected from April 2008 to October 2010. Both datasets share the same format: ⟨user ID, latitude, longitude, timestamp, location ID⟩. Each dataset also contains a social network of friendships, which serves as the ground truth in our evaluation.[2] Some statistics of these two datasets are given in Table I.

In this paper, we conduct both quantitative evaluations and case studies to verify the effectiveness of our method. For

[2]However, it is worth noting that the cyber (i.e., online) relationship may not serve as the *perfect* ground truth for the physical (i.e., mobility) relationship. We will discuss some examples using case studies in Section VI-D.

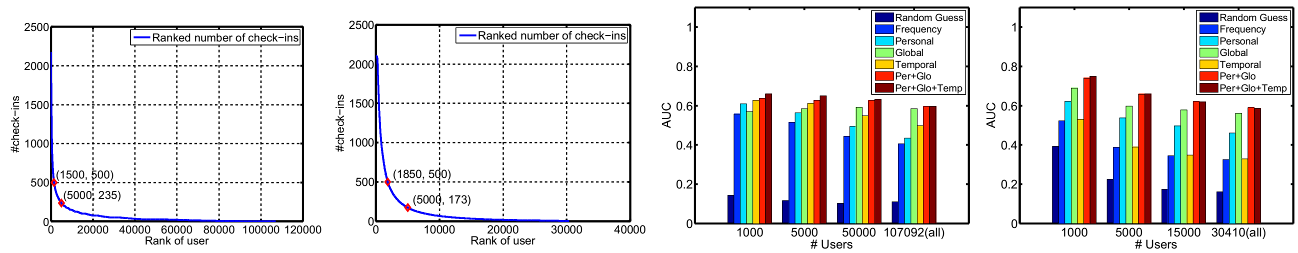

(a) #check-ins, Gowalla (b) #check-ins, Brightkite (c) AUC values, Gowalla (d) AUC values, Brightkite

Fig. 3. Effectiveness of our method. **(a) and (b)** show the ranked number of check-ins for all users in Gowalla and Brightkite datasets, respectively. It shows a long tail distribution. **(c) and (d)** show the AUC values of all methods for top-K users w.r.t. various values of K. The overall performance will decrease with more inactive users. All the factors and the combined factors can greatly outperform the baseline frequency measure.

TABLE I
DATASET STATISTICS

	Gowalla	Brightkite
No. of users	107,092	58,228
No. of friend pairs	950,327	214,078
No. of check-ins	6,442,890	4,491,143
Average check-ins per user	60	78

the former, we use the *precision-recall curve* to systematically examine each component in our framework and make comparison with the state-of-the-art method. Let G denote the set of ground truth friend pairs in the dataset, and Q denote the set of friend pairs reported by any method under a particular experiment setting, the precision and recall of the method are then defined as follows:

$$Precision = \frac{|G \cap Q|}{|Q|}, \quad Recall = \frac{|G \cap Q|}{|G|}. \tag{16}$$

To further compare different methods under various experiment settings, we adopt the *area under the curve* (AUC) value, which is the area under a given precision-recall curve. Finally, in order to show how to select the optimal PGT value as cutoff threshold, we also employ the F1 measure, which is

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}. \tag{17}$$

B. Effectiveness of Our Method

In this section, we systematically study the performance of each factor and the combinations as well. The following variations are considered:

- **Frequency:** $|E_{ij}|$ (meeting frequency).
- **Personal:** $\max_{e_k \in E_{ij}}\{w_{ij}^p(e_k)\} \times |E_{ij}|$ (meeting frequency weighted by the personal factor only).
- **Global:** $\sum_{e_k \in E_{ij}} w_{ij}^g(e_k)$ (meeting frequency weighted by the global factor only).
- **Temporal:** $\sum_{e_k \in E_{ij}} w_{ij}^t$ (meeting frequency weighted by the temporal correlation only).
- **Personal + Global:** $\max_{e_k}\{w_{ij}^p(e_k)\} \times \sum_{e_k} w_{ij}^g(e_k)$ (meeting frequency weighted by both personal factor and global factor).
- **Personal + Global + Temporal:** $\max_{e_k}\{w_{ij}^p(e_k)\} \times \sum_{e_k}\left(w_{ij}^g(e_k) \times w_{ij}^t(e_k)\right)$ (the PGT model which considers all the factors).

To gain insights into the datasets, in Figure 3(a) and 3(b) we plot the ranked number of user check-ins in both datasets. As shown, the number of check-ins for each user is highly skewed and both distributions have long tails. In particular, only $1,500$ out of the $100,000$ users in the Gowalla dataset have more than 500 check-ins. Similar phenomenon can be seen in the Brightkite dataset. This indicates that *a small portion of users use the services very frequently, but most of the users are quite inactive*. So one important question here is how the performance of our method is affected by the different degrees of user activeness.

To answer this question, we extract the top-K users (ranked by the number of check-ins) in the dataset and test our methods on all user pairs among these top-K users which co-locate at least once. In Figure 3(c) and 3(d), we show the AUC values of all methods w.r.t. different values of K. From the results, we can make the following observations.

1. *The meeting frequency itself is an important indicator of relationships.* It achieves much higher accuracies than random guess, which shows the ratio of the friend pairs in both datasets. Further, it works well for users with large number of check-ins (i.e., top-ranked active users), but its performance degrades as the degree of user activeness decreases. In other words, it is more difficult to predict the relationships of inactive users using the meeting frequency alone.

2. *Both the personal factor and the global factor help better differentiate the friend pairs from non-friend pairs.* But personal factor is more sensitive to the number of check-ins of each user, since its performance degrades faster than the global factor as the value of K increases. This is because we rely on each user's own location history to compute his personal background. So the more check-ins a user has, the more accurate the estimation of his personal background is. Finally, by combining the personal factor and the global factor, we are able to leverage the strengths of both factors to achieve a consistently better performance over the baseline.

3. *The performance of our method may be further improved by considering the temporal dependencies between consecutive events.* In particular, the temporal correlation weight works well for users with large number of check-ins, as it effectively penalizes the consecutive meeting events between two users with a small time gap. However, it becomes less useful as the level of user activities decreases (so that more user pairs

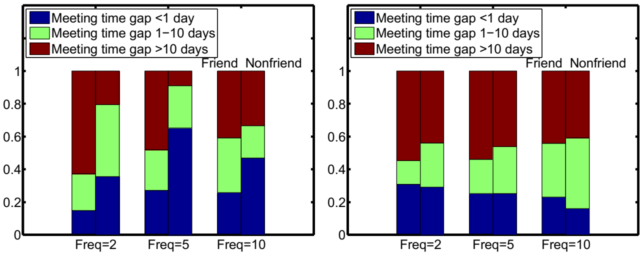

(a) Gowalla (b) Brightkite

Fig. 4. The distribution of time gaps between consecutive meeting events for three representative groups (meeting frequency $= 2, 5, 10$). For each group, we plot the percentages of time gaps (Δ_t) of all friend pairs (left) and non-friend pairs (right) that fall into one of the three categories: $\Delta_t < 1$, $1 \leq \Delta_t < 10$, and $\Delta_t \geq 10$.

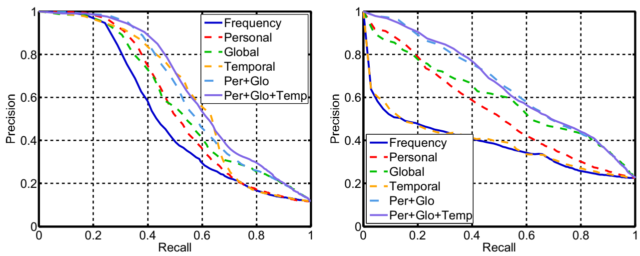

(a) Gowalla (b) Brightkite

Fig. 5. The precision-recall curves on top 5000 users from the two datasets.

with one or two co-locating events are included). Indeed, in the extreme case when the meeting frequency of a user pair is 1, the temporal dependency factor reduces to the baseline meeting frequency measure.

In addition, comparing Figure 3(c) with 3(d), we can see that the temporal correlation weight is more effective in the Gowalla dataset than the Brightkite dataset (also see Figure 5 for more details). To understand this phenomenon, we need to examine some inherent properties of the two datasets. In Figure 4, we show the distributions of time gaps of consecutive meeting events in the two datasets. As shown in Figure 4(a), in Gowalla dataset, the average time gap of friend pairs is significantly larger than that of non-friend pairs. For example, among all user pairs with meeting frequency of 5, more than 50% of time gaps for friend pairs are larger than 10 days, whereas more than 50% of the time gaps for non-friend pairs are less than 1 day. In fact, we note that many co-locating events of non-friend pairs occur in a very short time span at popular places such as Times Square. Thus, by considering the temporal correlations, we can effectively differentiate them from those of friend pairs, which tend to occur in a more regular basis over a long time span. However, as shown in Figure 4(b), in Brightkite dataset there is no significant difference in terms of the distribution of time gaps for friend and non-friend pairs. As a result, the temporal correlation measure is not very useful.

In Figure 5(a) and (b), we further plot the precision-recall curves of all methods for top 5000 users in Gowalla and Brightkite datasets, respectively. Looking closely at the results, we can see that the personal factor is more effective when recall is low (i.e., for pairs with higher meeting frequency,

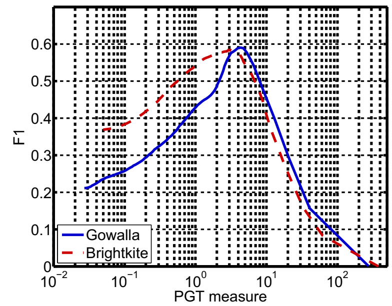

Fig. 6. The F1 curves for two dataset using different $\texttt{PGT}$ threshold to identify friend pairs.

because they are more easily to be retrieved), whereas the global factor is more effective when recall is high (i.e., for the pairs with lower meeting frequency). This can be explained as follows: many non-friend pairs co-locate at popular public places by chance, but friend pairs tend to meet at less popular places. For a pair with one or two co-locating events, using global factor one can easily differentiate whether they are friends or not. Meanwhile, friends with higher meeting frequency (i.e., stronger relationships) are more likely to travel together to some interesting places (e.g., a national park), where they rarely visit by themselves. Such cases can be easily captured by the personal background model.

In order to show how to pick the optimal threshold for $\texttt{PGT}$, we plot the F1 scores w.r.t. different $\texttt{PGT}$ thresholds in Figure 6. We treat all the pairs with $\texttt{PGT}$ values higher than the threshold as friend pairs. For Gowalla and Brightkite, our method achieves the best performance at $(F1 = 0.59, \texttt{PGT} = 4.28)$ and $(F1 = 0.58, \texttt{PGT} = 2.95)$, respectively. When setting $\texttt{PGT}$ threshold in the range of $[3, 4]$, the F1 scores are higher than 0.57 on both datasets. This implies that $\texttt{PGT}$ measure is consistent on the two datasets, in spite of their differences, such as the difference of time gaps shown in Figure 4.

C. Parameter Sensitivity

In our methods, there are two parameters, c_d and c_t, which control how fast the impact of a single location record decays in spatial and temporal dimensions, respectively. In this section, we study the sensitivity of our method to these two parameters using the Gowalla dataset.

Spatial parameter: c_d is used to estimate the density of a specific location for our personal factor. Note that, in this paper, we use kilometer as the unit of spatial distance. In Figure 7(a), we show the precision of our personal background measure ($G_2(E_{ij})$) as a function of c_d at different recall levels (0.3, 0.5, and 0.7). As one can see, this measure achieves the best performance when c_d is in the range of $[1, 3]$. Note that if c_d is in this range, the exponential function gets relatively large values when $d < 1km$, and quickly approaches zero when $d > 5km$. This is expected, considering the typical range of human movements.

Temporal parameter: c_t is used to calculate the temporal dependencies among multiple meeting events of a user pair. Note that, in this paper, we use day as the unit of time. In

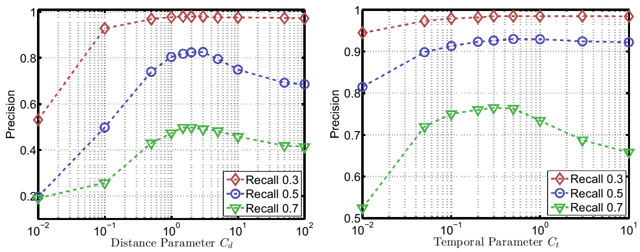

(a) (a) (b) (b)

Fig. 7. Parameter sensitivity. **(a)** Precision of $G_2(E_{ij})$ w.r.t. c_d. **(b)** Precision of $G(E_{ij})$ w.r.t. c_t. Note that our personal factor weight is positively correlated to $e^{-c_d \cdot d}$, while our temporal weight is positively correlated to $1 - e^{-c_t \cdot t}$.

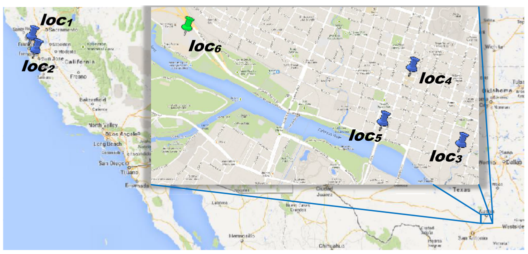

Fig. 8. The meeting places of two user pairs. Both pairs have met five times in total. The first (friend) pair meet twice in San Francisco (loc_1 and loc_2), and three times in downtown Austin (loc_3 to loc_5,). The second (non-friend) pair meet five times at the same location loc_6.

Figure 7(b), we plot the precision of our relationship strength measure ($G(E_{ij})$) as a function of c_t at different recall levels. We can see that our measure works the best when c_t is in the range of $[0.1, 0.3]$. This result suggests that, if two meeting events occur within about one week (7 days), they are less independent, and therefore should be down-weighted when we compute the relationship strength.

Based on the above analysis, we set $c_d = 1.5$ and $c_t = 0.2$ for all the experiments in this paper.

D. Case Studies

In this section, we perform case studies aiming to gain insights into why our proposed method works and when it fails to work on some cases.

Case 1 (Background factors): In this case, we study two user pairs with the same meeting frequency and see how our background factors can be used to differentiate them. The first pair of users (#267 and #510) are friends, whereas the second pair of users (#350 and #6138) are not friends. Both pairs have met for five times in total.

In Figure 8, we show all the meeting locations for both pairs on the map. In particular, users #267 and #510 (the friend pair) meet at five different locations with two in San Francisco (loc_1 and loc_2) and three in downtown Austin (loc_3, loc_4, and loc_5). Meanwhile, users #350 and #6138 (the non-friend pair) meet five times all in the same place loc_6.

TABLE II
LOCATION ENTROPY OF THE 6 LOCATIONS

loc_1	loc_2	loc_3	loc_4	loc_5	loc_6
1.609	3.184	4.775	1.282	3.998	2.189

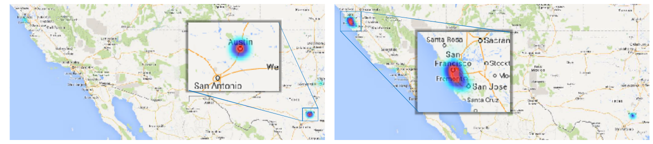

(a) User #267 (b) User #510

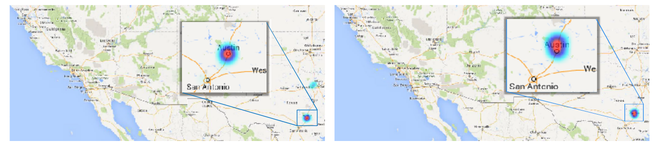

(c) User #350 (d) User #6138

Fig. 9. The density map of visited locations for the four users. Majority of user #510's activities are in San Francisco, whereas the other three users mainly live in Austin.

To examine the global factor, we further show the location entropies of these locations in the Table II. Compared to locations loc_2, loc_3 and loc_5, location loc_6 has a relatively low entropy. If we compute the average global factor weights for these two pairs, we get 1.52 for users #267 and #510, and 1.57 for users #350 and #6138. Therefore, *if we simply look at the global factor, we cannot get confident conclusion. We may even reach the wrong conclusion that users #350 and #6138 are more likely to be friends*.

Looking into the personal factor, in Figure 9(a)-(d), we show the density map of the locations visited by each of the four users, respectively. We can easily see that the majority of user #510's activities happen in San Francisco, whereas all the other three users mainly live in Austin. Given the fact that users #267 and #510 live in two different cities, the probability that they meet each other five times all by coincidence is very low. So they are very likely to be friends and have travelled to see each other intentionally. In contrast, both users #350 and #6138 live in Austin and frequently visit loc_6 so they are more likely to meet by chance (location loc_6 is actually a high school and these two users could attend the same high school but do not personally know each other). In this case, personal factor plays an important role in differentiating the relationship. If we calculate the frequency weighted by the personal factor, we get 22.03 for #267 and #510 (the friend pair) and 9.72 for #350 and #6138 (the non-friend pair).

Case 2 (Temporal correlation): In this case study, we add another non-friend pair of users (#39746 and #39584) who also meet five times. In Table III, we compare the personal factor weight ($\max_{e_k}\{w^p_{ij}(e_k)\}$), the global factor weight ($\overline{w}^g_{ij}$) and the temporal correlation weight ($\overline{w}^t_{ij}$) of all the three user pairs. As one can see, the personal factor weight for the this pair (i.e., the third pair in Table III) is 23.80, which is higher than the other two pairs. This suggests that they meet at places where they rarely visit, contradicting to our intuition that non-friend pairs tend to co-locate at places where they frequently visit.

However, if we look closely at their meeting locations, we can see that these locations are quite far away from each other. In particular, the distance from loc_8 to loc_{10} is 857km, but the time difference between these two events is only about 8 hours. Therefore we can infer that #39746 and #39584 must

TABLE III
THE THREE MEASURES FOR THREE PAIRS

User ID		Friends?	Personal	Global	Temporal
267	510	Yes	22.03	1.52	3.77
350	6138	No	9.72	1.57	3.99
39746	39584	No	23.80	2.23	1.39

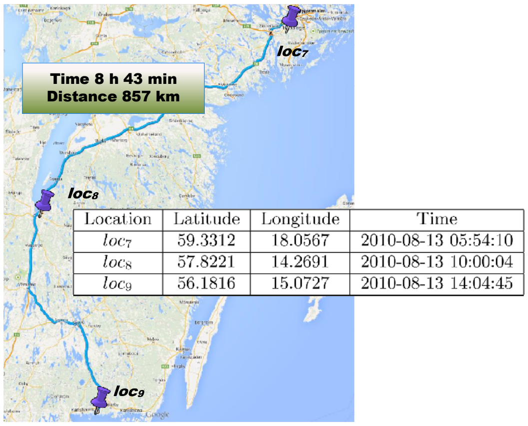

Fig. 10. Three meeting events happened within 8 hours for a non-friend pair #39746 and #39587.

be traveling together (possibly on a train), and have made three consecutive check-ins in the same day. Such meeting events can be captured using the temporal correlation weight, as shown in Table III. Combining all three factors, users #267 and #510 still have the highest overall score among the three pairs, which agrees with the fact that they are the only friend pair in this case.

Case 3 (Failed Cases): It is also interesting to study some failed cases, such as those non-friend pairs which are assigned with high relationship scores by our method. For example, the highest ranked non-friend pair in the Gowalla dataset is users #10683 and #10681. Looking at their location history, we find that they meet 24 times in total, which is higher than 90.97% of the friend pairs. In addition, the 24 meeting events occur at 16 different locations, and are spread out over one-year period. All these facts indicate that they should be friends. So we further look into the social networks of users #10683 and #10681. We find that they share 16 common friends, while they have 51 and 36 friends, respectively. Therefore, a possible explanation for this abnormal case may be that they simply did not report their friendship using the social networking service.

We also check the friend pairs that are given low mobility relationship scores by our method. In particular, the lowest ranked friend pair is users #6248 and #4609, who only meet once at the popular Vimeo Theater in downtown Austin. The global factor will suggest they are non-friend. These facts suggest that they should not be friends, and the reason for this abnormal case may be two-fold: (1) the mobility data is too sparse to estimate their true mobility relationship, and (2) they could have some online interactions which are not captured in their mobility data.

E. Comparison with the State-of-the-Art Method

In this experiment, we further compare our `PGT` with EBM [14], a state-of-the-art method which achieves the best

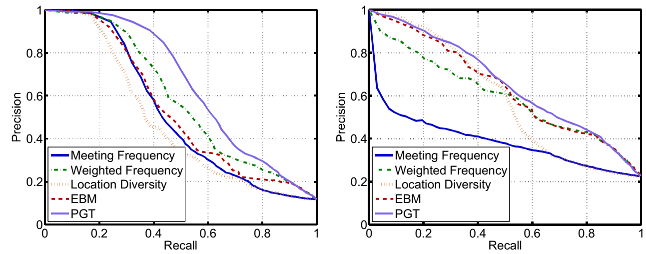

(a) Gowalla (b) Brightkite

Fig. 11. Comparison with the state-of-the-art methods.

performance among previous works [17], [15], [13], [18]. EBM is an entropy-based model designed to infer social strength from the users' movement data. It considers the following two major factors:

Location Diversity: Two users who meet at many different places are more likely to be socially connected than users who only meet at one or two places, even if both pairs have the same meeting frequency. In EBM, the Renyi entropy is used to measure the diversity of co-occurrences. In general, the more places two users have met, the higher the diversity will be.

Weighted Frequency: EBM also uses the location entropy to measure the popularity of a place. A popular public place visited by many users will have a higher location entropy than a private place. Using the same formula as our global factor, a weight is computed for each meeting event based on its location entropy.

Finally, a linear regression model is used in [14] to combine these two factors into the EBM model (in this paper, we use the same coefficients as reported in [14]).

In Figure 11, we show the precision-recall curves of various methods including (1) the meeting frequency (baseline), (2) the weighted frequency (meeting frequency weighted by the global factor/location entropy), (3) the location diversity, (4) EBM (linear combination of (2) and (3)), and (5) our social strength measure (Personal + Global + Temporal). From the results, we can make the following observations.

1. The location diversity measure alone performs poorly on the Gowalla dataset (worse than the meeting frequency measure), but works relatively well on the Brightkite dataset. To explain this inconsistency, we show the average number of co-locating places for various groups of users in both datasets in Table IV. As one can see, friend pairs in Brightkite dataset indeed visit more diverse places together than non-friend pairs. But this is not true for the Gowalla dataset.

TABLE IV
AVERAGE NUMBER OF CO-LOCATING PLACES.

Data set		$freq = 2$	$freq = 5$	$freq = 10$
Gowalla	Friend Pair	1.85	4.03	7.26
	Non-friend Pair	1.81	4.36	7.12
Brightkite	Friend Pair	1.53	2.92	4.47
	Non-friend Pair	1.19	1.44	1.18

2. As expected, the weighted frequency outperforms the meeting frequency on both datasets, since it considers the global background of each place. In addition, since EBM combines

weighted frequency and location diversity using a linear model combining weighted frequency and location diversity, it does not work well on Gowalla dataset due to the poor performance of the location diversity measure. But EBM outperforms both the weighted frequency and location diversity components on the Brightkite dataset.
3. Our method performs the best among all. In addition to the global background considered in EBM, our method also takes into account the personal mobility background and the temporal correlations among multiple events. As shown in our experiments, each of these factors play an important role in differentiating actual meeting events between friends and other co-locating events between strangers.

VII. Related Work

Using the geographical records to infer people's social behaviors and relationships is a hot topic in spatiotemporal data mining. Extensive research has been done in this area.

One related area of research is on the similarity measure of trajectories [19], [20], [21], [22]. Here, the subject of study is not restricted to sequences of locations on the map, but also include trajectories in other spaces, such as the positions of body joints of a person playing Kung-Fu, hand-writing trajectories, or hurricane trajectories. *Such measures are not suitable for judging the similarity of human movements.*

To measure of similarity of human trajectories, the *sequence similarity* has been studied in the literature [23], [24], [25]. In these works, a human trajectory is first transformed into a sequence of semantic locations, such as "shopping mall $\rightarrow$ restaurant $\rightarrow$ cinema". The similarity between two trajectories is then measured as a weighted matching score between the symbolic sequences. Such measures can be used to find people who have taken the same routes in their movement history, but *they may not necessarily appear in the same place at the same time.*

A co-locating event captures the direct interaction among moving objects. Based on the co-locating events, a line of research has been focused on mining moving clusters from the spatiotemporal data. Representative works include [8], [9], [10], [11], [12]. However, all these methods measure the degree of relationships based on the meeting frequency, and do not consider the background models of the moving objects. As we have shown in this paper, *the meeting frequency itself may not necessarily indicate the actual relationships.*

Our work falls in the category of studies that aim to detect social relationship from geospatial data [17], [15], [13], [18], [14]. The methods proposed in [17], [15] have looked into the meeting events that occur at different times (e.g., weekday v.s. weekend or day v.s. night) to infer different types of relationships such as colleagues and friends. Meanwhile, Cranshaw et al. [15] extract a set of features from both the meeting events and the individual mobility patterns and learn a model to identify friendships in social check-in data. In addition, Pham et al. [18], [14] further consider the diversity of meeting locations to handle cases that two users meet by coincidence. However, *none of these methods has considered using the personal background to differentiate meeting events.*

VIII. Conclusion

In this paper, we have studied the problem of measuring the relationship strength of mobile users based on their spatiotemporal interactions. We have proposed a unified framework to integrate different types of background models, together with the temporal correlation of multiple meeting events. Extensive experiments on two real datasets show that our method significantly outperforms the state-of-the-art in discovering true relationships in the mobile users.

References

[1] L. Liu, A. Hou, A. Biderman, C. Ratti, and J. Chen, "Understanding individual and collective mobility patterns from smart card records: A case study in shenzhen," in *Proc. ITSC*, 2009.
[2] E. L. Glaeser and M. E. Kahn, "Sprawl and urban growth," in *Cities and Geography*, ser. Handbook of Regional and Urban Economics, 2004.
[3] Q. Hao, R. Cai, C. Wang, R. Xiao, J.-M. Yang, Y. Pang, and L. Zhang, "Equip tourists with knowledge mined from travelogues," in *Proc. WWW*, 2010, pp. 401–410.
[4] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang, "Collaborative location and activity recommendations with gps history data," in *Proc. WWW*. ACM, 2010, pp. 1029–1038.
[5] S. Dhar and U. Varshney, "Challenges and business models for mobile location-based services and advertising," *Commun. ACM*, 2011.
[6] K. Shin, X. Ju, Z. Chen, and X. Hu, "Privacy protection for users of location-based services," *Wireless Communications, IEEE*, 2012.
[7] Y. Ge, H. Xiong, C. Liu, and Z.-H. Zhou, "A taxi driving fraud detection system," in *Proc. ICDM*, 2011.
[8] P. Kalnis, N. Mamoulis, and S. Bakiras, "On discovering moving clusters in spatio-temporal data," in *Proc. SSTD*, 2005, pp. 364–381.
[9] G. Al-Naymat, S. Chawla, and J. Gudmundsson, "Dimensionality reduction for long duration and complex spatio-temporal queries," in *Proc. SAC'07*, 2007.
[10] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen, "Discovery of convoys in trajectory databases," in *Proc. VLDB*, 2008.
[11] Z. Li, B. Ding, J. Han, and R. Kays, "Swarm: Mining relaxed temporal moving object clusters," in *Proc. VLDB*, 2010.
[12] K. Zheng, Y. Zheng, N. J. Yuan, and S. Shang, "On discovery of gathering patterns from trajectories," in *Proc. ICDE*, 2013.
[13] D. J. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg, "Inferring social ties from geographic coincidences," *PNAS*, vol. 107, no. 52, pp. 22 436–22 441, 2010.
[14] H. Pham, C. Shahabi, and Y. Liu, "Ebm: An entropy-based model to infer social strength from spatiotemporal data," in *Proc. SIGMOD*, 2013.
[15] J. Cranshaw, E. Toch, J. I. Hong, A. Kittur, and N. Sadeh, "Bridging the gap between physical location and online social networks," in *Proc. Ubicomp*, 2010.
[16] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: user movement in location-based social networks," in *Proc. KDD*, 2011.
[17] N. Eagle, A. Pentland, and D. Lazer, "Inferring friendship network structure by using mobile phone data," in *PNAS*, 2009.
[18] H. Pham, L. Hu, and C. Shahabi, "Towards integrating real-world spatiotemporal data with social networks," in *Proc. GIS*, 2011.
[19] B.-K. Yi, H. V. Jagadish, and C. Faloutsos, "Efficient retrieval of similar time sequences under time warping," in *Proc. ICDE*, 1998, pp. 201–208.
[20] M. Vlachos, D. Gunopulos, and G. Kollios, "Discovering similar multidimensional trajectories," in *Proc. ICDE*, 2002, pp. 673–684.
[21] L. Chen, M. T. Özsu, and V. Oria, "Robust and fast similarity search for moving object trajectories," in *Proc. SIGMOD*, 2005.
[22] L. Chen and R. T. Ng, "On the marriage of lp-norms and edit distance," in *Proc. VLDB*, 2004.
[23] Q. Li, Y. Zheng, X. Xie, Y. Chen, W. Liu, and W.-Y. Ma, "Mining user similarity based on location history," in *Proc. GIS*. ACM, 2008.
[24] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, "Mining interesting locations and travel sequences from gps trajectories," in *Proc. WWW*, 2009.
[25] X. Xiao, Y. Zheng, Q. Luo, and X. Xie, "Finding similar users using category-based location history," in *Proc. GIS*, 2010, pp. 442–445.