

Quantifying Location Sociality

Jun Pang
FSTC & SnT
University of Luxembourg
jun.pang@uni.lu

Yang Zhang
CISPA, Saarland University
Saarland Informatics Campus
yang.zhang@uni-saarland.de

ABSTRACT

The emergence of location-based social networks provides an unprecedented chance to study the interaction between human mobility and social relations. This work is a step towards quantifying whether a location is suitable for conducting social activities, and the notion is named location sociality. Being able to quantify location sociality creates practical opportunities such as urban planning and location recommendation. To quantify a location's sociality, we propose a mixture model of HITS and PageRank on a heterogeneous network linking users and locations. By exploiting millions of check-in data generated by Instagram users in New York and Los Angeles, we investigate the relation between location sociality and several location properties, including location categories, rating and popularity. We further perform two case studies, i.e., friendship prediction and location recommendation, experimental results demonstrate the usefulness of our quantification.

CCS CONCEPTS

• Information systems → Social networking sites; • Human-centered computing → Ubiquitous and mobile computing;

KEYWORDS

Online social networks; location-based social networks; data mining; friendship prediction; location recommendation

ACM Reference format:

Jun Pang and Yang Zhang. 2017. Quantifying Location Sociality. In *Proceedings of HT '17, Prague, Czech Republic, July 04-07, 2017*, 10 pages. DOI: <http://dx.doi.org/10.1145/3078714.3078729>

1 INTRODUCTION

Online social networks (OSNs) have been the most successful web applications during the past decade. Leading companies, including Facebook,¹ Twitter² and Instagram,³ have gained a large number of users. More recently, with the development of positioning technology on mobile devices, OSNs have been extended to geographical space. Nowadays, it is quite common for OSN users to share their geographical locations, i.e., check-ins. Moreover, a special type of

¹<https://www.facebook.com/>

²<https://twitter.com/>

³<https://www.instagram.com/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HT '17, Prague, Czech Republic

© 2017 ACM. 978-1-4503-4708-2/17/07...\$15.00

DOI: <http://dx.doi.org/10.1145/3078714.3078729>

OSNs dedicated to location sharing are created, namely location-based social networks (LBSNs). Foursquare and Yelp are two representative companies. With the emergence of LBSNs, a large quantity of data concerning human mobility become available. This gives us an unprecedented opportunity to understand human mobility and moreover to study the interaction between social relations and mobility. Some previous works have been focused on inferring social relationships from mobility, such as [13, 28, 32, 37, 41], others exploit users' social information to predict their future locations, such as [1, 3]. More recently, researchers propose new understandings of locations by using user generated data such as happiness [30] and walkability [29].

Location has been recognized as an important factor for social activities back in 1950s. In his seminal work [10], Erving Goffman described social interactions as a series of performance given by social actors, and physical setting, i.e., location, is an important aspect of a social actor's performance. In [10], Goffman stated that "A setting tends to stay put, geographically speaking, so that those who would use a particular setting as a part of their performance cannot begin their act until they have brought themselves to the appropriate place". Based on Goffmann's study, Milligan [21] further proposed that "physical sites (however defined by the participants) become the stages for social interaction, stages that are both physically and socially constructed". She explained that not only being physically constructed (by architects, facility managers, property owners and others), a location will also be socially constructed by people who conduct social interactions there. Following this theory, we argue that social construction will make some locations more suitable for social activities than others.

In the current work, we aim to quantify whether a location is a suitable for conducting social activities. The notion we quantify is named *location sociality*. We define a location's sociality as *the degree to which individuals tend to conduct social activities at that location*. A location is considered social if friends frequently visit, especially for the purpose of socializing or recreation, and vice versa. Studying location sociality could advance the boundary of our understanding on the interaction between social relations and mobility. It can also help us to solve challenging problems such as urban planning and traffic control. In practice, location sociality can be also used to build appealing applications such as location recommendation.

Contributions. In the current work, we make the following contributions:

- We propose a framework to quantify location sociality (Section 2). Our framework is based on the assumption that a location's sociality and its visitors' social influence are mutually reinforced. To model this assumption, we construct a heterogeneous network consisting of users (in a social

network) and locations (a user-location network). Then, we propose a mixture model of HITS [16] and PageRank to quantify location sociality on this heterogeneous network.

- Following our solution, we exploit millions of check-in data from Instagram in New York and Los Angeles to quantify location sociality (Section 3). We then study the relation between location sociality and several location properties including location categories, rating and popularity. Our discoveries include: certain types of locations (music venues and nightclubs) are more social than others; location sociality shares a positive relation with location rating given by users; social locations distribute more uniformly w.r.t. geographical space than popular locations.
- To demonstrate the usefulness of our quantification of location sociality, we perform a case study on friendship prediction in Section 4.1. We extract two users' common locations and define features based on these common locations' sociality for machine learning classification. Experimental results show that with very simple location sociality features, we are able to achieve a strong prediction. Moreover, adding location sociality into a state-of-the-art prediction model achieves a 5% performance gain.
- We perform another case study on using location sociality for location recommendation in Section 4.2. We integrate our quantification into a random walk with restart framework. Experimental results show that the recommender based on location sociality achieves a better recommendation performance (at least 5%) than the baseline recommender that does not consider location sociality.

We discuss some implications and limitations of the current work in Section 5. Related works are discussed in Section 6 and Section 7 concludes the paper.

2 PROPOSED SOLUTION

In this section, we first discuss the intuition of our solution on quantifying location sociality (Section 2.1), then we formally describe the solution (Section 2.2).

2.1 Intuition

Our intuition on quantifying location sociality in this paper is based on the assumption that a location's sociality and its visitors' social influence are mutually reinforced. To explain this intuition, we start by addressing socially influential users. In the society, if a person is considered socially influential, he must visit different social places frequently to organize or participate in different social activities and events. On the other hand, if a location is frequently visited by influential users, then it must be suitable for conducting social activities, i.e., it is a social place. Following this, we establish a mutual reinforcement relation between user influence and location sociality, i.e., more social a location is, more socially influential users visit it, and vice versa. In addition to visiting many social places, an influential user should also occupy an important position in the social network, e.g., he should have many friends who are also socially influential. Following the above discussion, our intuition on quantifying location sociality can be summarized as the following two assumptions.




Figure 1: A model of the network.

Assumption 1. Location sociality and users' social influence are mutually reinforced.

Assumption 2. Users' social influence can be quantified from the social network.

The first intuition can be naturally formulated into a HITS-style framework [16]. For the second intuition, we apply PageRank on the social graph to quantify each user's influence. In the rest of the paper, we use location sociality and sociality interchangeably.

2.2 Our framework

We start by modeling users, locations and their relationships into two types of networks including social network and user-location network.

Social network. A *social network*, denoted as $G_U = (U, \mathcal{E}_U)$, is an unweighted graph with nodes in set U representing all users. $\mathcal{E}_U \subseteq U \times U$ is a symmetric relation containing the edges in G_U . If u_i and u_j are friends, then $(u_i, u_j) \in \mathcal{E}_U$ and $(u_j, u_i) \in \mathcal{E}_U$. We use matrix X to represent G_U where $X_{i,j} = 1$ if $(u_i, u_j) \in \mathcal{E}_U$ and $X_{i,j} = 0$ otherwise. It is easy to see that X is symmetric. We further use \bar{X} to denote the column stochastic matrix of X where $\bar{X}_{i,j} = \frac{X_{i,j}}{\sum_k X_{k,j}}$.

User-location network. A *user-location network*, denoted as $G_{U,\mathcal{L}} = (U, \mathcal{L}, \mathcal{E}_{U,\mathcal{L}})$, is a weighted bipartite graph. $\mathcal{E}_{U,\mathcal{L}} \subseteq U \times \mathcal{L}$ consists of the edges in $G_{U,\mathcal{L}}$. Each edge $(u_i, \ell_j) \in \mathcal{E}_{U,\mathcal{L}}$, also written as $e_{i,j}^{U,\mathcal{L}}$, is associated with a weight $w_{i,j}^{U,\mathcal{L}}$ defined as the number of times that the user u_i has visited (checked in) the location ℓ_j (denoted by $|ci(u_i, \ell_j)|$). We use matrix Y to represent $G_{U,\mathcal{L}}$ with $Y_{i,j} = w_{i,j}^{U,\mathcal{L}}$. The transpose of Y is further denoted by Y^T . In the end, we use \bar{Y} and \bar{Y}^T to denote the column stochastic matrices of Y and Y^T , respectively.

Figure 1 shows an example of the heterogeneous graph. Within our framework two sets of values, locations' sociality and users'

social influence, can be obtained. Each location ℓ 's sociality is defined as $\kappa(\ell)$ and $\eta(u)$ for each user's social influence. Following the intuition in Section 2.1, our model is formulated into the following equations:

$$\eta(u_i) = \sum_j \bar{X}_{i,j} \cdot \eta(u_j) \quad (1)$$

$$\eta(u_i) = \sum_j \bar{Y}_{i,j} \cdot \kappa(\ell_j) \quad (2)$$

$$\kappa(\ell_j) = \sum_i \bar{Y}_{j,i}^T \cdot \eta(u_i) \quad (3)$$

Equations 1 is the PageRank implementation for quantifying users' social influence from $G_{\mathcal{U}}$. Equations 2 and 3 are an instance of the HITS framework which establishes the mutual reinforcement relationship between locations and users. We then linearly combine the above equations as

$$\eta(u_i) = \alpha \cdot \sum_j \bar{X}_{i,j} \cdot \eta(u_j) + (1-\alpha) \cdot \sum_j \bar{Y}_{i,j} \cdot \kappa(\ell_j) \quad (4)$$

$$\kappa(\ell_j) = \sum_i \bar{Y}_{j,i}^T \cdot \eta(u_i) \quad (5)$$

where α specifies the contributions of each component to users' social influence. In our experiments, α is set to 0.5 which indicates the social network structure and user mobility are equally important on quantifying users' social influence. Note that $\alpha = 0.5$ is a typical setting in many fields such as [36] where the authors aim to discover salient sentences for document summarization.

We further use two vectors η and κ to denote users' social influence and locations' sociality. Then the above equations can be written into the following matrix form.

$$\eta = \alpha \cdot \bar{X} \cdot \eta + (1-\alpha) \cdot \bar{Y} \cdot \kappa \quad (6)$$

$$\kappa = \bar{Y}^T \cdot \eta \quad (7)$$

Equations 6 and 7 can be computed through an iterative updating process. We set all locations' (users') initial sociality (social influence) to be $\frac{1}{|\mathcal{L}|}$ ($\frac{1}{|\mathcal{U}|}$). According to our experiments, the computation stops after around 10 iterations, when the maximal difference between κ s of two consecutive iterations is less than 0.00001.

3 EXPERIMENTS

In this section, we first introduce the dataset used for our experiments. Then, we present the results of our quantification: we start by discussing the top social locations and location categories; then we focus on the relation between location sociality and location rating; in the end, the correlation between location sociality and popularity is discussed.

3.1 Dataset description

Instagram is a photo-sharing social network with a fast growing user number. By now, it has 400M monthly active users and with 75M photos published everyday. Similar to other social network services such as Facebook and Twitter, Instagram allows users to share their locations when publishing photos. Moreover, unlike Twitter where only a small amount of tweets are geo-tagged, the authors of [19] have shown that Instagram users are much more willing to share their locations (31 times more than Twitter users),




Figure 2: Check-ins in New York.

	New York	Los Angeles
# check-ins	6,181,169	4,705,079
# active users	12,280	8,643
# edges (active users)	74,230	44,994
# locations	8,683	6,908

Table 1: Dataset summary.

which makes Instagram a suitable platform to study the interaction between mobility and social relations.

We collect the geo-tagged photos, i.e., check-ins, in New York and Los Angeles from Instagram through its public API⁴. Since locations' category information is an important aspect of our analysis, and fortunately the API of Instagram is linked with the API of Foursquare, a leading location-based social network with resourceful information about each place, thus we exploit the following methodology to collect our data. We first resort to Foursquare to extract all location ids within each city, meanwhile we collect each location's category information together with its rating (number of tips and number of likes). Then for each Foursquare's location id, we query Instagram's API to get its corresponding location id in Instagram. After this, we query each location's recent check-ins in Instagram several times a day from August 1st, 2015 until March 15th, 2016. In the end, more than 6M check-ins have been collected in New York and 4.7M in Los Angeles⁵. To resolve the data sparseness issue, we focus on users with at least 20 check-ins (considered as active users) and locations with at least 10 check-ins. Figure 2 depicts a sample check-in distribution in New York. Since Foursquare organizes location categories into a tree structure⁶, we take its second level categories to label each location.

To obtain users' social networks, we exploit Instagram's API to query each active user's follower/followee list⁷. We consider two users as friends if they mutually follow each other in Instagram. To further guarantee that users we have collected are not celebrities or business accounts, we filter out the top 5% of users with most

⁴<https://www.instagram.com/developer/>

⁵It is worth noticing that the authors of [20] has applied a similar methodology.

⁶<https://developer.foursquare.com/categorytree>

⁷Since Instagram's API only provides one page with 50 follower/followees per query, we perform multiple queries until all follower/followees of each user are obtained.

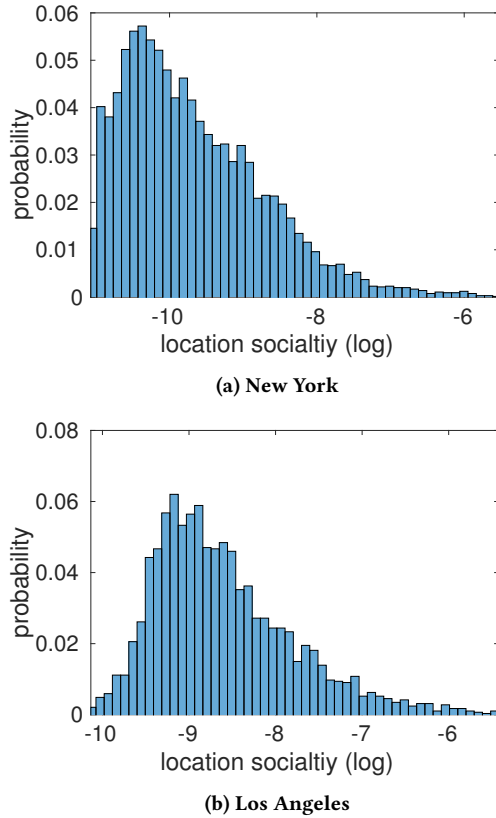


Figure 3: Distributions of log-transformed location sociality.

followers. Also, only the relations among active users (users with at least 20 check-ins) are kept. In the end, the social network contains 74,230 edges for New York and 44,994 edges for Los Angeles. Table 1 summarizes the dataset. For the sake of experimental result reproducibility, our dataset is available upon request.

3.2 Location sociality vs. location category

Figure 3 depicts the log transformed distributions of location sociality, both of which indicate that most locations have a middle value of sociality while only a few locations are very social or unsocial. This is different from other location measurement, for instance, the number of mobility transitions from or to each location follows a power law distribution [24].

Table 2 presents the top five locations with highest and lowest sociality. For New York, Webster Hall (music venue) is the most social place followed by Madison Square Park (park). On the other hand, the least social place is one Staples store (convenience store) in midtown. For Los Angeles, The Fonda Theatre (concert hall) has the highest sociality. Meanwhile, one Panda Express (fast food restaurant) is the least social place. From Table 2, we can see a clear distinction between social and unsocial places w.r.t. their categories. Next, we take a deeper look at the relation between sociality and location category.

New York	
Top Social	Top Unsocial
Webster Hall	Staples
Madison Square Park	17 Frost Gallery
Rockwood Music Hall	China Institute
Washington Square Park	Manhattan Theatre Club
Baby's All Right	El Rey Del Taco II
Los Angeles	
Top Social	Top Unsocial
The Fonda Theatre	Panda Express
Avalon Hollywood	Gap
The Echo	Ebar
Hermosa Beach Pier	Palms Super Market
Exchange LA	7-Eleven

Table 2: The most and least social locations in New York and Los Angeles.

New York	
Social Categories	Unsocial Categories
Music Venue	Laundry Service
Nightclub	Convenience Store
Harbor	Post Office
Museum	Pharmacy
Park	Fast Food Restaurant
Los Angeles	
Social Categories	Unsocial Categories
Concert Hall	Convenience Store
Nightclub	Vintage Store
Music Venue	Fast Food Restaurant
Mall	Pet Service
Beach	Automotive Shop

Table 3: Top 5 location categories with highest and lowest average location sociality in New York and Los Angeles.

Table 3 lists the top five location categories with the highest and lowest average location sociality. Nightclub and music venue are in the top 3 in both cities. On the other hand, convenience store seems to be less attractive to friends. Besides, we also observe some interesting difference between the two cities. For example, beach is the No.5 social choice for people living in Los Angeles while it is not New Yorkers' choice since there are no beaches in Manhattan.

As music venue and nightclub have high rankings in both cities, we further list the top 5 music venues and nightclubs in Table 4. Although the ranking of music venues and nightclubs are rather subjective, we have checked several blogs and articles (listed in the additional material) and most of our top-social nightclubs and music venues have received positive reviews and been recommended by these blogs and articles. We conclude that a location's sociality

New York	
Social Music Venues	Social Nightclubs
Webster Hall	Stage 48
Rockwood Music Hall	Marquee
Baby's All Right	Pacha NYC
Bowery Ballroom	1 OAK
Music Hall of Williamsburg	VIP Room NYC
Los Angeles	
Social Music Venues	Social Nightclubs
Avalon Hollywood	Exchange LA
The Echo	OHM Nightclub
The Roxy	Sound Nightclub
The Troubadour	Club Los Globos
The Hollywood Bowl	Create Nightclubs

Table 4: Top 5 music venues and nightclubs with highest location sociality in New York and Los Angeles.

is related to its category. Normally, location category itself is not sufficient to judge whether a location is social or not. Next, we study other properties of locations and their relationship with location sociality.

3.3 Location sociality vs. rating, tips and likes

For each location, Foursquare provides us with not only its category information, but also other properties including rating⁸, number of tips and number of likes generated by Foursquare users. Next, we study whether it is possible to use these properties to explain location sociality. To proceed, we build a linear regression model with rating, number of tips and number of likes as explanatory variables while location sociality as the dependent variable. By fitting the model with ordinary least square method, we obtain a coefficient of determination (R^2) of 0.192 in New York and 0.280 in Los Angeles, meaning that 19.2% (28.0%) of the variability of location sociality in New York (Los Angeles) can be explained by these properties. By checking the parameters of our linear model, we discover that the major predictive power is driven by location rating.

We further plot the average location sociality as a function of rating in Figure 4: the two variables share a positive relation. Especially when location rating is high (≥ 8), location sociality increases sharply for both cities. This indicates that social places are assigned with high ratings by users.

3.4 Location sociality vs. location popularity

A social location is often popular in the sense that it attracts many people. On the other hand, to conduct social activities, everyone has his own preference on choosing locations. For example, one may prefer to go to a bar near his home, which might not be well-known at the city level. The relationship between a location's sociality and its popularity is worth investigation: a location's sociality should be

⁸In Foursquare, rating is in the range from 1 to 10.

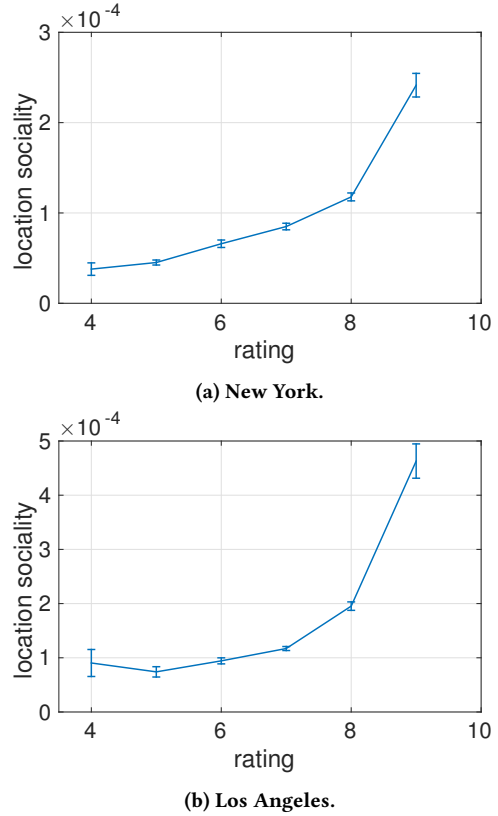


Figure 4: Location sociality vs. location rating.

correlated with its popularity, while the two notions should exhibit difference.

By far, the most common notion for quantifying a location's popularity is location entropy [4], it is formally defined as

$$le(\ell) = - \sum \frac{|ci(u, \ell)|}{|ci(\ell)|} \log \frac{|ci(u, \ell)|}{|ci(\ell)|},$$

where $|ci(u, \ell)|$ is user u 's number of check-ins at location ℓ (Section 2) and $|ci(\ell)|$ is the total number of check-ins of location ℓ . More popular a location is, higher location entropy it has.

To check the difference between the two measurements, we plot the heatmaps w.r.t. location entropy and sociality of New York. As expected, midtown and downtown New York are "hot" areas in both maps. On the other hand, we observe that location sociality is more uniformly distributed than location entropy. For example, the areas marked by green circles in Figure 5a are obviously lighter than those in Figure 5b. After having a close look, we discover that bars and restaurants are the "hot" locations inside these areas. Data in Los Angeles exhibits a similar result and is not shown.

We further extract the top 20 popular locations in both cities (listed in the additional material). In New York, the most popular locations are parks and museums (e.g., The MET, MoMA and Guggenheim). On the other hand, in Los Angeles, the most popular locations concentrate on malls followed by museums. Moreover, in both cities, famous landmarks have high location entropy, such as Rockefeller Center in New York and Hollywood Walk of Fame in




Figure 5: Heatmaps in New York.

Los Angeles. This is quite different from the ranking in Table 2 and Table 3: social locations are mainly music venues and nightclubs while popular locations are mainly tourist attractions. In the end, we conclude that there exists a large difference between social and popular locations.

4 CASE STUDIES

In this section, we perform two case studies to show how location sociality can be used to build real world applications. The applications we focus on are friendship prediction and new location recommendation, both of which are essential for social network services.

4.1 Friendship prediction

Following the seminal work of Liben-Nowell and Kleinberg, friendship prediction has been extensively studied [17], resulting in appealing applications such as friendship recommendation which is essential for OSNs to increase user engagement. During the past five years, with the development of LBSNs, many researchers start to exploit users' location data as a new source of information for friendship prediction.

Model. In our case study, we consider friendship prediction as a binary classification problem. Each pair of friends is treated positive if they are friends (mutually following each other in Instagram) and negative otherwise. We extract the common locations of two users and construct the feature space based on these common locations.

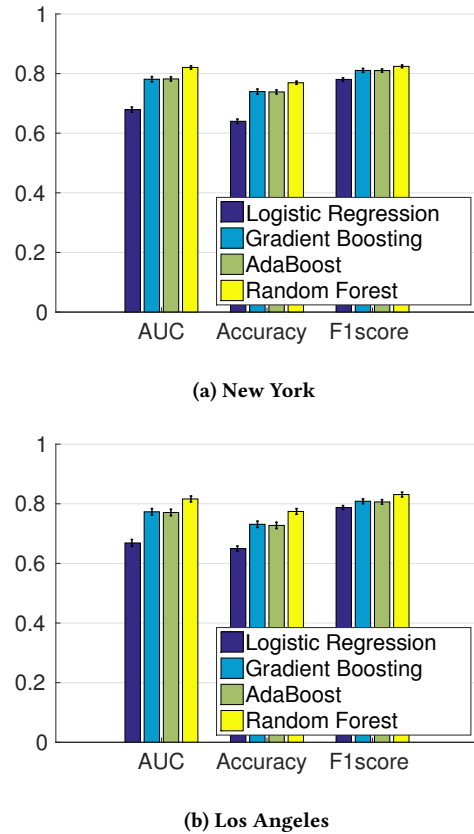


Figure 6: Evaluation results on our model w.r.t. four classification algorithms.

Here, two users' common locations are the intersection of the places they have checked in, regardless of time. For two users u_i and u_j , we find their common locations' sociality and utilize the average, maximal, minimal and standard deviation of these sociality as features for classification.

Experiment setup. To resolve the data sparseness issue, we filter out pairs of users who have only one or zero common location. This leaves us 7,525 pairs of friends, i.e., positive cases, in New York and 3,961 pairs of friends in Los Angeles. For negative case, we randomly sample the same number of non-friend pairs – each pair of them has at least two common places as well. It is worth noticing that this way of sampling negative cases increases the hardness of classification since a non-friend pair also has at least two common locations. Therefore, we can further evaluate the usefulness of location sociality. We have adopted four classification algorithms in our experiments including logistic regression, gradient boosting, AdaBoost and random forest. Accuracy, F1score and AUC (area under the ROC curve) are used as our metrics. We randomly split the dataset with 70% for training and 30% for testing, this random split is repeated for 10 times and we report the average result.

Results. Figure 6 depicts the performance of our classifications. Among all the classifiers, random forest performs the best with

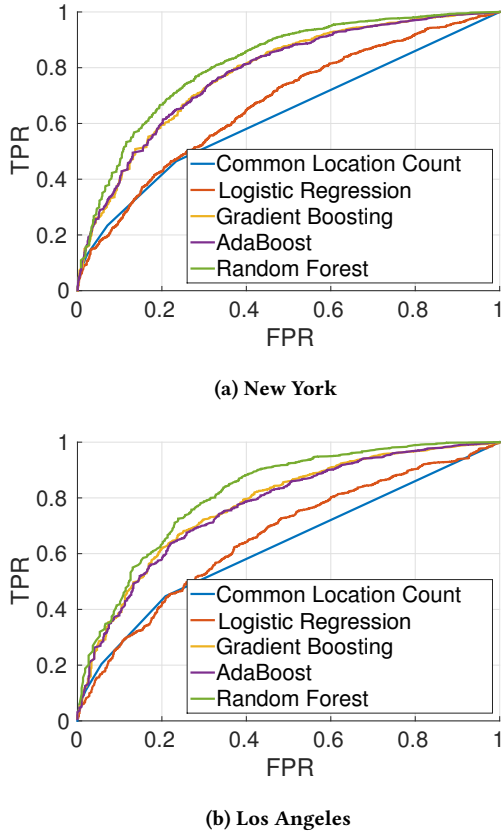


Figure 7: ROC curves for our classification and the number of common locations.

AUC = 0.82 and Accuracy = 0.77 in the two cities. Meanwhile, we have F1score = 0.82 in New York and F1score = 0.83 in Los Angeles. AdaBoost and gradient boosting have a comparable performance. On the other hand, logistic regression performs the worst.

The number of common locations is further adopted as a naive baseline model for comparison, i.e., we tune the threshold (the number of common locations) for classification to obtain the ROC curve. Figure 7 plots the results. As we can see, all our classifiers based on location sociality outperform this naive baseline.

Next, we check whether adding location sociality into a state-of-the-art model as proposed in [32]⁹ can increase prediction performance. The model in [32] (location feature setting) extracts two users' common locations and design features mainly with these common locations' entropies (see Section 3.4 for location entropy), such as the minimal location entropy. In our experiments, we combine our four location sociality features with the features in [32] and fit them into our best performing classifier random forest. The results in Table 5 show that the classification with location sociality improves [32] by around 5% among all three metrics in both

⁹Most of the recent solutions focus on two users' meeting events i.e., two users at the same location at roughly the same time, while [32] and our model only exploit common locations.

New York	AUC	Accuracy	F1score
[32]	0.83	0.77	0.82
sociality+[32]	0.87	0.81	0.85
Los Angeles	AUC	Accuracy	F1score
[32]	0.82	0.77	0.82
sociality+[32]	0.86	0.80	0.85

Table 5: Evaluation result on [32] and location sociality+[32].

New York and Los Angeles. This further demonstrates that location sociality is useful for friendship prediction.

Indeed, there exist other solutions for friendship prediction such as considering two users' meeting events [13, 28, 37]. However, the main issue for this method is that meeting events (reflected in OSNs) are rare, for example, with 6M check-ins in New York, we only observe around 100 meeting events. Moreover, we want to emphasize that friendship prediction is not the focus of the current work. Therefore, we choose the most straightforward features for training our classifiers. Nevertheless, our prediction still achieves a strong performance showing that location sociality is a good indicator for recommending friends. Further investigation on integrating location sociality into the state-of-the-art friendship prediction models is worth studying and we leave it as a future work.

4.2 Location recommendation

The second case study we perform is recommending new locations for users to visit. Location recommendation has a great potential to build appealing applications. During the past five years, it has attracted academia a lot of attention (e.g. [2, 6, 7, 42]). Our goal here is to demonstrate the usefulness of location sociality in recommending new locations. In order to integrate location sociality into a location recommender, we adopt a classical approach, namely random walk with restart [34].

Model. In a typical setting of random walk with restart for recommendation, in the beginning we define a matrix Q as

$$Q = \begin{pmatrix} 0 & Y \\ Y^T & 0 \end{pmatrix}$$

where Y and Y^T represent user-location network (location-user network) (Section 2). Meanwhile, \bar{Q} denotes the column stochastic version of Q [39]. Then to recommend locations to a user u_i , we modify \bar{Q} to allow every node in the graph having a certain probability (15% in the experiments) to jump to the node representing u_i . Formally, for every $Q_{a,b} \in Q$, $\bar{Q}_{a,b}$ is defined as

$$\bar{Q}_{a,b} = \begin{cases} (1-c) \cdot \frac{Q_{a,b}}{\sum_j Q_{j,b}} + c \cdot 1 & \text{if the } a\text{th row represents } u_i \\ (1-c) \cdot \frac{Q_{a,b}}{\sum_j Q_{j,b}} & \text{otherwise} \end{cases}$$

where $c = 0.15$. By applying the same method of solving PageRank, e.g., power method, we can obtain the steady state distribution over \bar{Q} , which is the relevance score of all nodes (both locations and users) to u_i . Locations with high relevance scores are recommended to u_i . Noulas et al. [23] have exploited this approach for location

recommendation¹⁰, where the weight on an edge between a user u_i and a location ℓ_j is simply the user's number of visits to that location, i.e., $Y_{i,j} = w_{i,j}^{\mathcal{U},\mathcal{L}} = |ci(u_i, \ell_j)|$ in Section 2.

To integrate location sociality into the edge weight for location recommendation, we change Y to T , i.e., Q is modified to:

$$Q = \begin{pmatrix} 0 & T \\ \mathcal{T} & 0 \end{pmatrix}$$

where $T_{i,j}$ is defined as

$$T_{i,j} = |ci(u_i, \ell_j)| \cdot \frac{1}{-\log(\kappa(\ell_j))}. \quad (8)$$

Here, $\kappa(\ell_j)$ is the location sociality of ℓ_j , meanwhile \mathcal{T} is the transpose of T . Under this formulation, Equation 8 assigns higher weight to locations with high sociality, which will bias the recommended locations to be more social. In the end, by performing power method on the column stochastic version of the modified Q , we obtain the recommended locations for each user.

Experiment setup. The check-in dataset is partitioned temporally with each one covers consecutively 60 days [23]. For each partition, we use the data of the first 30 days to train the model while the left 30 days for testing. Since our aim is to perform new location recommendation, for each user we further filter out his locations in the testing set that he has already been to in the training set. In the end, we perform random walk with restart with location sociality ($rwr-1s$) to recommend locations for each user, and exploit rwr without location sociality (rwr), i.e., the one in [23], as the baseline model. Two metrics including precision@10 and recall@10 are adopted for evaluation.

Results. Table 6 presents the results for location recommendation in both cities. As we can see, $rwr-1s$ outperforms rwr in all months. For precision@10, $rwr-1s$ outperforms rwr by 10%, while for recall@10, even in the worst case in New York, $rwr-1s$ still has 3.4% improvement on rwr . Even though the absolute precision and recall of our recommendation is not high, it is worth noticing that the similar performances of location recommendation have been obtained by [7, 18, 38], thus our results are reasonable. Similar to Gao et al. [7], we emphasize that the focus here is to compare the relative performance, in order to demonstrate the usefulness of our quantification.

Many state-of-the-art algorithms exploit other factors for location recommendation such as geographical distance and users' published contents, one of our future works is to integrate location sociality into these algorithms to further improve recommendation.

5 IMPLICATIONS AND LIMITATIONS

We discuss implications and limitations of the current work.

Implications. Location sociality as a measurement can characterize a social map for a city. The applications based on it, including the friendship prediction and location recommendation addressed in Section 4, can benefit several parties.

For city administrator, understanding where people prefer to socialize can help them make better city plans. For example, during conventional social time such as Friday nights, the city government

could make specific transportation plans, such as more buses and taxis, or security plans such as deploying more policemen, for areas with high sociality. Moreover, location sociality can be a good reference for the government to plan future city development. For city residents, location sociality provides a good reference for them to find or discover new places to socialize with their friends. In addition, location sociality can be used as an important factor for location recommendation services for social network services, such as Yelp and Foursquare. For visitors, visiting high sociality places is a good way to engage local people's social life. This can help visitors understand the city's culture in a better way. For business owners, knowing where people like to go to conduct social activities is an important factor for them to determine where to open new business.

Limitations. We point out the following limitations of the current work. First, we only focus on the data from Instagram which cannot reflect the general population both socially and geographically. Socially, the authors of [22, 33] have shown that most of Instagram users are center around a younger age (25 years old); geographically, most of the check-ins concentrate on the city center (Figure 2).

Second, our quantification does not take into account the temporal factor which could also be important for understanding locations. One approach for considering time would be concentrating on users' meeting events, i.e., users check in at the same location at the same time, however, as we have discussed previously, meeting events are rare even in large datasets as ours. In the future, we plan to incorporate data with rich temporal information from other sources to further improve our quantification.

Third, we quantify a location's sociality based on its visitors' information. In some cases, a location's own property can also contribute to its sociality. For example, it is pointed out that the decoration, the space allocation and even the bicycle parking design can result in different number of visitors to a cafe located in a university campus [21]. However, we argue that data of this kind is hard and expensive to obtain at a large scale.

6 RELATED WORK

The emergence of LBSNs has brought us an unprecedented opportunity to study human mobility and its interaction with social relationships. Many works have been done on understanding human mobility and its interaction with social relations. There mainly exist two research directions: one is to exploit users' location information to understand social relations, e.g., see [13, 28, 32, 37, 41]; the other is to use social relations to understand mobility and locations, including the current work.

Backstrom et al. [1] present one of the pioneer works on friendship-based location prediction. They analyze Facebook users' home location and discover that friends tend to live closer to each other than strangers. Then they build a maximal likelihood estimator to predict a user's home location. They have shown that their model outperforms significantly the method based on IP addresses. Cho et al. [3] study a general problem: instead of predicting home location, they aim to predict where a user is at a certain time. They construct a dynamic Gaussian mixture model with the assumption that each user's mobility is centered around two states, such as home and work. The experimental results show that their model

¹⁰They [23] also consider social network in Q , here we ignore it for better demonstrating location sociality's usefulness.

New York						
15.8-15.10	Precision@10	Recall@10		15.11-16.1	Precision@10	Recall@10
rwr	0.009	0.021		rwr	0.009	0.028
rwr-ls	0.010	0.024		rwr-ls	0.010	0.031
15.9-2015.11	Precision@10	Recall@10		15.12-16.2	Precision@10	Recall@10
rwr	0.010	0.032		rwr	0.009	0.026
rwr-ls	0.011	0.034		rwr-ls	0.010	0.028
15.10-15.12	Precision@10	Recall@10		16.1-16.3	Precision@10	Recall@10
rwr	0.009	0.028		rwr	0.008	0.027
rwr-ls	0.010	0.029		rwr-ls	0.009	0.029
Los Angeles						
15.8-15.10	Precision@10	Recall@10		15.11-16.1	Precision@10	Recall@10
rwr	0.010	0.028		rwr	0.008	0.024
rwr-ls	0.011	0.030		rwr-ls	0.009	0.026
15.9-2015.11	Precision@10	Recall@10		15.12-16.2	Precision@10	Recall@10
rwr	0.013	0.038		rwr	0.012	0.047
rwr-ls	0.015	0.042		rwr-ls	0.013	0.048
15.10-15.12	Precision@10	Recall@10		16.1-16.3	Precision@10	Recall@10
rwr	0.010	0.025		rwr	0.007	0.035
rwr-ls	0.012	0.029		rwr-ls	0.009	0.044

Table 6: Precision@10 and recall@10 for location recommendation.

achieves a promising accuracy. Other works include [8, 14, 26, 27]. More recently, Jurgens et al. [15] perform a comprehensive study on most of the existing works in the field and points out some future directions.

Besides predicting a user’s location, researchers begin to advance our understandings of locations based on the data from social networks. In [30], the authors focus on recommending pleasant paths between two locations in a city. Unlike the traditional shortest path recommendation, they assign three values to describe whether a street is quiet, beautiful and happy, respectively. Then they adjust the path recommendation algorithm with these factors and recommend the most pleasant path for users. In [29], the authors quantify whether a street is suitable for walk, namely walkability. To assess their results, they propose to use concurrent validity. Their discoveries, to mention a few, include walkable streets tend to be tagged with walk-related words on Flickr and can be identified by location types on those streets. The authors of [5] exploit the data from Foursquare to analyze different neighborhoods in a city. They extract some signature features to profile each neighborhood and propose an algorithm to match similar neighborhoods across different cities. Experimental results show that they are able to match tourists areas across Paris and Barcelona, and expensive residential areas in Washington D.C. and New York. More recently, Hristova et al. [11] propose four location measurements, including brokerage, serendipity, entropy and homogeneity, under a heterogenous social and location network model. Their experiments are conducted with a Foursquare dataset collected in London, and the authors show

that their proposed measurements can be used to describe dynamics that is hard to capture including gentrification and deprivation. Other recent works include [9, 12, 25, 31].

The current work also falls into the field of urban informatics [43], a newly emerging field where researchers tend to use the ubiquitous data to understand and improve the city where we live. Besides the research literature, several open projects have been established as well. To mention a few examples, Yuan et al. [40] focus on discovering the function of each region in a city; Venerandi et al. [35] measure the socio-economic deprivation of a city. Another excellent example is the goodcitylife project¹¹, where the team members try to imitate human beings’ five senses on food to understand cities.

7 CONCLUSION

In this paper, we have proposed a new notion namely location sociality to describe whether a location is suitable for conducting social activities. We constructed a heterogenous network linking locations and users and proposed a mixture model of HITS and PageRank to quantify location sociality. Experimental results on millions of Instagram check-in data validate location sociality with some in-depth discoveries. Two case studies including friendship prediction and location recommendation demonstrate the usefulness of our quantification.

¹¹<http://goodcitylife.org/>

Location data do not only come from LBSNs, but many other sources, such as GPS traces and WIFI points. In the future, we are interested in establishing more connections between LBSN data and other sources to gain a deep understanding of cities.

REFERENCES

- [1] Lars Backstrom, Eric Sun, and Cameron Marlow. 2010. Find me if you can: improving geographical prediction with social and spatial proximity. In *Proc. 19th International Conference on World Wide Web (WWW)*. ACM, 61–70.
- [2] Jie Bao, Yu Zheng, David Wilkie, and Mohamed Mokbel. 2015. Recommendations in location-based social networks: a survey. *Geoinformatica* 19, 3 (2015), 525–565.
- [3] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *Proc. 17th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 1082–1090.
- [4] Justin Cranshaw, Eran Toch, Jason Hone, Ankit Kittur, and Norma Sadeh. 2010. Bridging the gap between physical location and online social networks. In *Proc. 12th ACM International Conference on Ubiquitous Computing (UbiComp)*. ACM, 119–128.
- [5] Geraud Le Falher, Aristides Gionis, and Michael Mathioudakis. 2015. Where Is the soho of Rome? Measures and algorithms for finding similar neighborhoods in cities. In *Proc. 9th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press, 228–237.
- [6] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2013. Exploring temporal effects for location recommendation on location-based social networks. In *Proc. 7th ACM Conference on Recommender Systems (RecSys)*. ACM, 93–100.
- [7] Huiji Gao, Jiliang Tang, Xia Hu, and Huan Liu. 2015. Content-Aware Point of Interest Recommendation on Location-Based Social Networks. In *Proc. 29th AAAI Conference on Artificial Intelligence (AAAI)*. The AAAI Press, 1721–1727.
- [8] Huiji Gao, Jiliang Tang, and Huan Liu. 2012. Exploring social-historical ties on location-based social networks. In *Proc. 6th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press, 114–121.
- [9] Petko Georgiev, Anastasios Noulas, and Cecilia Mascolo. 2014. Where businesses thrive: predicting the impact of the Olympic games on local retailers through location-based services data. In *Proc. 8th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press, 151–160.
- [10] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Random House.
- [11] Desislava Hristova, Matthew J. Williams, Mirco Musolesi, Pietro Panzarasa, and Cecilia Mascolo. 2016. Measuring Urban Social Diversity Using Interconnected Geo-Social Networks. In *Proc. 25th International Conference on World Wide Web (WWW)*. ACM, 21–30.
- [12] Hsun-Ping Hsieh, Cheng-Te Li, and Shou-De Lin. 2015. Estimating Potential Customers Anywhere and Anytime Based on Location-Based Social Networks. In *Proc. 2015 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD)*. Springer, 576–592.
- [13] Hsun-Ping Hsieh, Rui Yan, and Cheng-Te Li. 2014. Where You Go Reveals Who You Know: Analyzing Social Ties from Millions of Footprints. In *Proc. 24th ACM International Conference on Information and Knowledge Management (CIKM)*. ACM, 1839–1862.
- [14] David Jurgens. 2013. That's what friends are for: Inferring location in online social media platforms based on social relationships. In *Proc. 7th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press.
- [15] David Jurgens, Tyler Finethy, James McCorriston, Yi Tian Xu, and Derek Ruths. 2015. Geolocation prediction in Twitter using social networks: a critical analysis and review of current practice. In *Proc. 9th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press, 188–197.
- [16] Jon Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (1999), 604–632.
- [17] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology* 58, 7 (2007), 1019–1031.
- [18] Bin Liu and Hui Xiong. 2013. Point-of-interest recommendation in location based social networks with topic and location awareness. In *Proc. 13th SIAM International Conference on Data Mining (SDM)*. SIAM, 396–404.
- [19] Lydia Manikonda, Yuheng Hu, and Subbarao Kambhampati. 2014. Analyzing user activities, demographics, social network structure and user-generated content on Instagram. *CoRR* abs/1410.8099 (2014).
- [20] Yelena Mejova, Hamed Haddadi, Anastasios Noulas, and Ingmar Weber. 2015. #FoodPorn: obesity patterns in culinary interactions. In *Proc. the 5th International Conference on Digital Health (DH)*. ACM, 51–58.
- [21] Melinda J Milligan. 1998. Interactional past and potential: The social construction of place attachment. *Symbolic Interaction* 21, 1 (1998), 1–33.
- [22] Minyue Ni, Yang Zhang, Weili Han, and Jun Pang. 2016. An Empirical Study on User Access Control in Online Social Networks. In *Proc. 21st ACM Symposium on Access Control Models and Technologies (SACMAT)*. ACM, 13–23.
- [23] Anastasios Noulas, Salvatore Scellato, Neal Lathia, and Cecilia Mascolo. 2012. A Random Walk around the City: New Venue Recommendation in Location-Based Social Networks. In *Proc. 2012 International Conference on Social Computing (SocialCom)*. IEEE, 144–153.
- [24] Anastasios Noulas, Blake Shaw, Renaud Lambiotte, and Cecilia Mascolo. 2015. Topological properties and temporal dynamics of place networks in urban environments. In *Proc. 24th International Conference on World Wide Web (WWW Companion)*. ACM, 431–441.
- [25] Alexandra-Mihaela Olteanu, Kévin Huguenin, Reza Shokri, Mathias Humbert, and Jean-Pierre Hubaux. 2017. Quantifying interdependent privacy risks with location data. *IEEE Transactions on Mobile Computing* 3 (2017), 829–842.
- [26] Jun Pang and Yang Zhang. 2015. Exploring communities for effective location prediction (poster paper). In *Proc. 24th International Conference on World Wide Web Conference (Companion Volume) (WWW)*. ACM, 87–88.
- [27] Jun Pang and Yang Zhang. 2015. Location prediction: Communities speak louder than friends. In *Proc. 3rd ACM Conference on Online Social Networks (COSN)*. ACM, 161–171.
- [28] Huy Pham, Cyrus Shahabi, and Yan Liu. 2013. EBM: an entropy-based model to infer social strength from spatiotemporal data. In *Proc. 2013 ACM International Conference on Management of Data (SIGMOD)*. ACM, 265–276.
- [29] Daniele Quercia, Luca Maria Aiello, Rossano Schifanella, and Adam Davies. 2015. The digital life of walkable streets. In *Proc. 24th International Conference on World Wide Web (WWW)*. ACM, 875–884.
- [30] Daniele Quercia, Rossano Schifanella, and Luca Maria Aiello. 2014. The shortest path to happiness: recommending beautiful, quiet, and happy routes in the city. In *Proc. 25th ACM Conference on Hypertext and Social Media (HT)*. ACM, 116–125.
- [31] Daniele Quercia, Rossano Schifanella, Luca Maria Aiello, and Kate McLean. 2015. Smelly maps: the digital life of urban smellscape. In *Proc. 9th AAAI Conference on Weblogs and Social Media (ICWSM)*. The AAAI Press, 237–236.
- [32] Salvatore Scellato, Anastasios Noulas, and Cecilia Mascolo. 2011. Exploiting place features in link prediction on location-based social networks. In *Proc. 17th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 1046–1054.
- [33] Flavio Souza, Diego de Las Casas, Vinicius Flores, SunBum Youn, Meeyoung Cha, Daniele Quercia, and Virgilio Almeida. 2015. Dawn of the selfie era: the whos, wheres, and hows of selfies on Instagram. In *Proc. 3rd ACM Conference on Online Social Networks (COSN)*. ACM, 221–231.
- [34] Hanghang Tong, Christos Faloutsos, and Jia yu Pan. 2006. Community detection in networks with node attributes. In *Proc. 6th IEEE International Conference on Data Mining (ICDM)*. IEEE CS, 613–622.
- [35] Alessandro Venerandi, Giovanni Quattrone, Licia Capra, Daniele Quercia, and Diego Saez-Trumper. 2015. Measuring urban deprivation from user generated content. In *Proc. 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*. ACM, 254–264.
- [36] Xiaojun Wan, Jianwu Yang, and Jianguo Xiao. 2007. Towards an iterative reinforcement approach for simultaneous document summarization and keyword extraction. In *Proc. 45th Annual Meeting of the Association of Computational Linguistics (ACL)*. 552–559.
- [37] Hongjian Wang, Zhenhui Li, and Wang-Chien Lee. 2014. PGT: Measuring mobility relationship using personal, global and temporal factors. In *Proc. 14th IEEE International Conference on Data Mining (ICDM)*. IEEE, 570–579.
- [38] Mao Ye, Xingjie Liu, and Wang-Chien Lee. 2012. Exploring social influence for recommendation: a generative model approach. In *Proc. 35th ACM SIGIR conference on Research and Development in Information Retrieval (SIGIR)*. ACM, 671–680.
- [39] Mao Ye, Dong Shou, Wang-Chien Lee, Peifeng Yin, and Krzysztof Janowicz. 2011. On the semantic annotation of places in location-based social networks. In *Proc. 17th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 520–528.
- [40] Jing Yuan, Yu Zheng, and Xing Xie. 2012. Discovering regions of different functions in a city using human mobility and POIs. In *Proc. 18th ACM Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 186–194.
- [41] Yang Zhang and Jun Pang. 2015. Distance and friendship: A distance-based model for link prediction in social networks. In *Proc. 17th Asia-Pacific Web Conference (APWeb)*. Springer, 55–66.
- [42] Vincent W. Zheng, Yu Zheng, Xing Xie, and Qiang Yang. 2010. Collaborative location and activity recommendations with GPS history data. In *Proc. 19th International Conference on World Wide Web (WWW)*. ACM, 1029–1038.
- [43] Yu Zheng, Licia Capra, Ouri Wolfson, and Hai Yang. 2014. Urban computing: concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology* 5, 3 (2014), 1–55.