

Exploring Venue Popularity in Foursquare

Yanhua Li*, Moritz Steiner[§], Limin Wang[§], Zhi-Li Zhang*, and Jie Bao*

*Department of Computer Science & Engineering, University of Minnesota, Twin Cities

[§]Bell Labs, Murray Hill, NJ

Email: {yanhua,zhzhang,baojie}@cs.umn.edu, {moritz,liminwang}@bell-labs.com

Abstract—In this paper, we provide a detailed analysis on the venue popularity in Foursquare, a leading location-based social network. By collecting 2.4 million venues from 14 geographic regions all over the world, we study the common characteristics of popular venues, and make the following observations. First, venues with more complete profile information are more likely to be popular. Second, venues in the Food category attract the most (43%) public tips (comments) by users, and the Travel & Transport category is the most popular category with the highest per venue check-ins, i.e., each venue in this category attracts on average 376 check-ins. Moreover, the stickiness of users checking in venues in the residence, office, and school categories is higher than in other categories. Last but not least, in general, old venues created at the early stage of Foursquare are more popular than new venues. Our results help to understand the factors that cause venues to become popular, and have applications in venue recommendations and advertisement in location based social networks.

I. INTRODUCTION

Location-based services have attracted remarkable interests over the last a few years, thanks to the fast deployment of broadband mobile networks and the increasing prevalence of versatile mobile devices. Leading online social networks, such as Facebook Places [2] and Google+ Local [12] have embedded location based services as an important feature. Foursquare [3], one of the most popular location based social networks (LBSNs) enables users to explore nearby places (e.g., for specials and discounts) and network with their friends. In September 2011, Foursquare had more than 10 million registered users with 1 billion check-ins [1], and by April 2012 the number of check-ins doubled [9].

The success of LBSNs is of great interest to those wishing to investigate and analyze large-scale data and their implications to improving location-based online services, e.g., user mobility prediction, friendship and venue recommendations [14], [15], [19]–[24].

As an important component in LBSNs, locations or venues play a crucial role in connecting users together and shedding light on understanding users’ preferences and mobility patterns. However, it is unclear why certain venues become popular, e.g., why they attract a large number of visits (check-ins) or comments (tips) from users, and what characteristics popular venues usually possess. Understanding and answering these questions is crucial to many applications, including venue recommendation and targeted advertising in LBSNs.

To answer these questions, in this paper, we make the first attempt to investigate how various factors affect the venue popularity in Foursquare, by analyzing a unique dataset consisting of 2.4 million venues collected from 14 geographic regions all over the world from May 1st to June 30th 2012. We first introduce some basics of Foursquare and LBSNs, and describe the data collection method and an overview of the dataset (Section II). Then we study several crucial characteristics that popularize the venues (Section III). Here is a summary of our main results:

- **Venue profile.** Venues with more complete profile are more likely to be popular, and the two most influential attributes are “contact” and “cross street” (Section III-A).
- **Venue category.** By performing comprehensive categorical analysis, we observe that venues in the Food category attract the most (43%) public comments (tips) by users, and the Travel & Transport category is the most popular category with the highest per venue check-ins, i.e., each venue in this category attracts on average 376 check-ins. Moreover, the residence, office, and school have the highest user stickiness (Section III-B).
- **Venue age.** The most popular venues were usually created at the early phase of Foursquare (Section III-C).

As a last contribution of this work, we made the anonymized venue dataset available [10] to the wider research community. We believe that our Foursquare data could facilitate a number of projects in social science and computer network research that try to look at characteristics of LBSNs.

II. BASICS OF FOURSQUARE AND THE VENUE DATASET

Foursquare [3] is a location-based social networking website launched in March 2009, and has become one of the most popular LBSNs. In this section, we briefly introduce the terminology used in this paper and provide an overview of our dataset, including the methodology we used for collecting the data and some statistics of the dataset.

A. Foursquare Venues

A Foursquare *venue* is a physical location. It can be a place of business, e.g., a restaurant, train station or movie theater, or a private residence. Foursquare users can create, check in, and leave tips to venues.

Creating a venue: Foursquare users can generate venues via Foursquare website or mobile applications. Each venue is assigned with a unique venue id, consisting of 24 hexadecimal characters. When creating a venue, a user is asked to provide a few attributes of the venue, such as the venue’s name, address,

This study was done partly while Yanhua Li was a summer intern at Bell Labs, Alcatel-Lucent. It was supported in part by the US NSF grant 0831734, CNS-1017647, the DTRA grant HDTRA1-09-1-0050, and a DoD ARO MURI Award W911NF-12-1-0385.

TABLE I. NUMBERS OF VENUES IN VARIOUS REGIONS (ORDERED BY CONTINENTS).

Regions	New York City	Los Angeles	Colorado	Orlando	Pittsburgh	Wyoming	Sydney
# non-home	453,070	193,262	166,190	46,945	29,312	14,179	12,721
Regions	Belgium	Switzerland	Paris	Slovenia	Cairo	Singapore	Seoul
# non-home	446,354	141,565	96,147	66,155	3,391	400,520	329,120

location, category, zip code, cross street, and etc. A *venue address* consists of a street number and a street name, and a *venue location* is specified by a user dropping a pin on the map in a Foursquare application. The application converts that location to a latitude and longitude tuple. Users choose the venue category from a list defined by Foursquare.

Foursquare defines a three-level hierarchical structure of categories for venues. There are nine top level categories: Arts & Entertainment, College & University, Food, Professional & Other Places, Nightlife Spot, Outdoors & Recreation, Shop & Service, Travel & Transport and Residence. Users need to specify a category when creating a venue. Foursquare launched the category feature on March 10th, 2010, so venues created before that date may not have any category information. A special case is the “Home (Private)” category, which is a subcategory of Residence. Venues in this category are privacy sensitive; Foursquare users can be understandably upset if they see strangers checking in at their homes. Hence, Foursquare has House Rules [8] suggesting that “Don’t check into someone else’s home if you’re not there” and “Only create a foursquare venue for someone else’s home if you have the permission of the resident/owner.” Moreover, Foursquare ensures that the sensitive details of a Home (Private) venue will be visible only to its owner and her friends. For example, other users will see a zoomed out view on the map instead of the venue’s precise location. In our collected dataset, the *home* venues’ location fields are fuzzed to an approximate area by Foursquare. Therefore, many “Home (Private)” venues in the same vicinity seem to have the same location.

Checking in a venue: Foursquare allows registered users to explicitly record their presence at a venue by making an active selection via its website or a mobile application. Check-ins can be either public or private. Users can choose to display their check-in information on their connected friends’ Foursquare sites, and post the check-ins on their Twitter or Facebook accounts. Check-ins will be awarded with points and coupons. Users can also choose to post their check-ins on their Twitter or Facebook accounts.

Tipping a venue: Foursquare users can add “Tips” to venues for other users to read. These tips serve as suggestions for great things to do, see, or eat at the location.

B. Venue Dataset

Foursquare’s search application programming interface (API) [4] can return a list of venues in a region. The region is specified by the latitudes and longitudes of the south-west and north-east corners of the region bounding box. The Foursquare API enforces a time constraint, i.e., 500 queries per hour per authenticated account, and a space constraint, i.e., up to 50 venues returned for a query. To ensure that we retrieve all the venues in a region, we keep the bounding box small enough so the number of venues in it does not exceed the return limit. This significantly limits the data collection speed. In our study,

we performed exhaustive search for two months in 2012 on 14 geographic regions using 40 machines, and collected in total 2,398,931 venues¹. The bounding box of each region is retrieved using Google GeoCoding API [11], e.g., the bounding box for New York City is $ne = (40.917577, -73.700272)$ and $sw = (40.477399, -74.25909)$. Table I lists the number of venues in each region. Our 14 regions cover a wide range of geographic areas, where the Foursquare services are popular [1]. As mentioned earlier, Foursquare fuzzes out the locations of Home (Private) venues to protect privacy, making these home venues information less accurate via the API. In addition, we primarily focus on analyzing venue popularity—the ability of venues attracting public users’ interest, while home venues tend to generate visits only from the owners themselves and their friends. Therefore, we only consider non-home venues in this study. Our anonymized venue dataset is available on [10].

Venue information. Each collected venue has two types of information: profile and statistics. A venue’s profile includes *user generated attributes*, such as venue name, address, cross street, latitude, longitude, city, state, country, zip code, contact (i.e., phone number), category, and the verification indicator, and *system generated attributes*, such as venue id and creation time. The verification indicator is a boolean variable, showing whether a business owner has already verified or claimed the venue, and a venue is by default unverified. A venue’s statistics include the total number of (public) check-ins, the total number of tips, and the total number of distinct users who have checked in the venue, since the venue has been created. The 2.4 million venues we have collected altogether generate over 308 million check-ins and 2 million tips¹. As observed in our dataset, the distributions of these three statistics for individual venues are Zipfian, which is consistent with the findings in many social networks [25].

These three statistics capture how popular a venue is, and in the next section, we will study how various aspects, such as the completeness of the venue profile information, venue category, and venue age, affect the venue popularity.

III. VENUE POPULARITY ANALYSIS

Foursquare users explicitly express their interests in venues in two ways, including checking in and leaving tips to venues. A *venue being frequently checked in* indicates that the venue is popular in a sense that lots of people visit it and like to announce their visits to their friends. A *venue being frequently tipped* indicates that people are interested in the venue and would like to share their experience with all other users. In the

¹We also circumvent the time cost of exhaustive data collection by designing an intelligent random region sampling algorithm, that allows us to collect venue samples and estimate the statistics, such as the total number of venues, and venue popularity distribution in a geographic region, with measurable variance (confidence interval). Please refer to [17] for more details.

¹Note that the number of users cannot be simply added up as the total number of distinct users, since a user might check in multiple venues.

TABLE II. AVAILABLE VENUE INFORMATION

Attributes	Name	Latitude	Longitude	Verif. indicator	Country	Category
Availability	100%	100%	100%	100%	99.99%	80.33%
Attributes	City	State	Address	Zip code	Contact	Cross street
Availability	62.37%	50.96%	48.33%	38.38%	13.60%	13.24%

TABLE III. INFLUENCE OF ATTRIBUTES ON VENUE POPULARITY

Attributes		Contact	Cross st.	Zip	Address	City	State	Category	Country
per venue # check-ins	with Attr.	565.45	485.66	267.53	238.49	189.30	205.37	153.45	128.53
	w/o Attr.	59.39	74.06	41.98	25.71	27.81	48.70	21.88	34.05
	influence	506.06	411.60	225.55	212.78	161.49	156.67	131.57	94.48
per venue # users	with Attr.	238.30	208.52	105.85	93.77	74.14	80.47	60.00	49.79
	w/o Attr.	19.97	25.59	14.89	8.67	9.45	17.92	6.13	19.76
	influence	218.33	182.93	90.96	85.10	64.69	62.55	53.87	30.03
per venue # tips	with Attr.	3.93	3.43	1.73	1.57	1.25	1.37	1.02	0.85
	w/o Attr.	0.37	0.46	0.31	0.18	0.20	0.32	0.16	0.79
	influence	3.56	2.97	1.42	1.39	1.05	1.05	0.86	0.06

following, we use the number of distinct users (who checked in the venue), the number of check-ins, and the number tips as three key statistics to analyze how the venue popularity is affected by three aspects, including the completeness of venue profile information, venue category, and venue age.

A. Effect of the completeness of venue profile information

We first investigate how the completeness of venue profile information impacts the venue popularity. When creating a venue in Foursquare, a user can provide a few attributes of the venue, such as the venue’s location, name, address, category, etc. The location, represented by a latitude and longitude tuple, is specified by the user dropping a pin on the map. However, these user generated attributes are mostly optional, except for the venue name, venue location, and verification indicator. Thus, many venues have missing attributes. Table II lists the availability for the attributes. Below, we will discuss in detail the influence of different attributes and the number of missing attributes on the venue popularity.

Influence of various user generated attributes. For each attribute that is not always present, we separate venues into two groups, i.e., with and without that specific attribute, and compute the average number of check-ins per venue for these two groups, respectively. The difference between the per venue check-ins of the two groups captures the attribute’s influence on the venue popularity in generating check-ins. The higher the difference is, the more influential the attribute is on the venue popularity. Similar analysis can be applied to the number of users and tips. Table III shows the venue popularity of each group with and without an attribute present, and the influence of each attribute, ordered by the degree of influence (decreasing from left to right). We observe that “contact” (e.g., phone number or Twitter ID) and “cross street” (i.e., the street crossed with the street in the address attribute) are the two most influential attributes. A possible explanation is that those attributes make the venue easier to reach, so venues with these attributes are more likely to be popular. Comparing with the Table II, it is interesting to note that these two attributes have the lowest availability ratios.

Effect of the number of missing attributes. We observe that a venue can have up to eight missing attributes, and we group venues by the number of missing attributes. Table IV shows the




Fig. 1. Effect of the completeness of venue profile information




Fig. 2. Popularity of verified vs unverified venues

venue distribution over the number of missing attributes. Note that most venues (about 25.43%) have six attributes missing. In Figure 1, we show the venue popularity in each group, measured by the ratio between the total number of check-ins (users and tips) to venues and the total number of venues in each group. The results clearly show that venues with fewer missing attributes are statistically more popular, i.e., they have more per venue check-ins, users, and tips.

TABLE IV. VENUE DISTRIBUTION OVER NUMBER OF MISSING ATTRIBUTES

# missing attri.	0	1	2	3	4	5	6	7	8
# total venues	101,084	239,323	348,242	378,023	302,001	237,321	610,146	182,422	369
% total venues	4.21%	9.98%	14.52%	15.76%	12.59%	9.89%	25.43%	7.60%	0.02%

TABLE V. VENUE DISTRIBUTION IN DIFFERENT CATEGORIES.

Index	Category Name	# venues	%	C5	Outdoors & Recreation	146,976	7.6%
C1	Arts & Entertainment	103,893	5.4%	C6	Professional & Other Places	434,890	22.6%
C2	College & University	87,973	4.6%	C7	Residence	65,370	3.4%
C3	Food	387,585	20.1%	C8	Shop & Service	398,867	20.7%
C4	Nightlife Spot	117,579	6.1%	C9	Travel & Transport	184,045	9.5%

Moreover, Foursquare allows business owners to verify or claim their venues, which is indicated by a true value of the verification indicator. However, among all 2,398,931 venues, there are only 73,939 of them (3%) verified. We next re-analyze how the number of missing attributes affect the popularity of the verified vs unverified venues, in terms of the per venue check-ins, users, and tips. Figure 2 clearly shows that verified venues have higher venue popularity than unverified venues.

B. Categorical analysis

In our dataset, there are 1,927,178 venues (around 80.33% of the total) with the category attribute available. Table V gives the numbers of venues in nine top level categories specified by Foursquare. Professional & Other Places (22.6%), Shop & Services (20.7%), and Food (20.1%) are the three largest categories. Under the nine top categories, there are in total 281 subcategories and 122 level-three categories (See a detailed list of all subcategories at [6]). Below, we will analyze and compare the venue popularity across different categories, in terms of the total and per venue check-ins and tips.

The total venue popularity distribution. Let C_i be one of the nine Foursquare venue categories, $i = 1, \dots, 9$, and $C(C_i)$ be the number of venue check-ins in category C_i . Let C be the total number of check-ins for all venues. The percentage of check-ins in each category, i.e., $Pr(C_i) = C(C_i)/C$, forms the check-in distribution among different categories. Similarly, we have the tip distribution representing the percentage of tips in different categories. (Again, we cannot do the same calculation on the number of users who has checked in a venue, since a single user may check in different venues in the same category and we cannot use simple addition to get the total number of distinct users for a particular category.) Figure 3 shows these two venue popularity distributions and the distribution of the total number of venues. We observe that the third largest category, Food (C3), consisting of 20.1% of total venues, generates far more tips than other categories, i.e., 43% of all tips were left to Food venues. On the other hand, Professional & Other Places (C6) and Shop & Services (C8) are the two largest categories, with 22.6% and 20.7% of the total number of venues, but they only attract 13–16% of check-ins, and 7–14% of tips. Moreover, the Travel & Transport category (C9) consisting of only 9.5% of the total venues attracts the most check-ins (around 23.4%).

Summary: These results imply that in absolute value, the Travel & Transport (C9) and the Food (C3) categories attract the most check-ins, whereas the Food (C3) category dominates other categories in generating users' interests of sharing tips.




Fig. 3. Venue and total venue popularity distribution over categories




Fig. 4. Normalized per venue popularity and user stickiness over categories

The per venue popularity analysis. We next consider the per venue popularity among different categories. For each category C_i , the average check-ins obtained by each venue, i.e., per venue check-ins, is counted as the total number of check-ins $C(C_i)$ divided by the total number of venues $n(C_i)$, i.e., $r(C_i) = C(C_i)/n(C_i)$, which reflects the average ability of venues in C_i to attract check-ins. Similarly, we have the per venue tips for each category. In addition, for a venue v , we define the user stickiness $s(v)$ as the ratio between the total number of check-ins to v and the total number of distinct users who checked in v . We use the average user stickiness of venues from C_i to evaluate the average ability of venues in C_i to keep frequent/recurrent visitors. Table VI

TABLE VI. RANKING OF PER VENUE POPULARITY

Rank	# of check-ins per venue		# of tips per venue		User stickiness	
1	Travel & Trans:	376	Food:	2.2	Residence:	7.83
2	Arts & Enter:	209	Travel & Trans:	1.4	Prof. & Others:	5.55
3	Outdoors & Rec:	160	Nightlife Spot:	1.3	College & U:	4.82
4	Nightlife Spot:	159	Arts & Enter:	1.2	Travel & Trans:	3.75
5	Food:	158	Outdoors & Rec:	0.7	Outdoors & Rec:	3.59
6	College & U:	126	Shop & Serv:	0.6	Nightlife Spot:	3.32
7	Shop & Serv:	121	College & U:	0.4	Arts & Enter:	3.28
8	Prof. & Others:	88	Prof. & Others:	0.4	Shop & Serv:	2.94
9	Residence:	62	Residence:	0.2	Food:	2.21

lists per venue popularity for each category in terms of per venue check-ins and tips, as well as user stickiness. Moreover, let $R = \sum_{i=1}^9 r(C_i)$, then the normalized per venue check-ins is computed as $\bar{r}(C_i) = r(C_i)/R$. Figure 4 shows the normalized per venue check-ins, tips, and the normalized user stickiness.

From Table VI and Figure 4, we observe that the Travel & Transport (C9) category generates the most per venue check-ins, which is 1.7 times higher than the average of the second highest category, Art & Entertainment (C1). In particular, Los Angeles International Airport (LAX) has the most check-ins in our dataset (i.e., 740,551 check-ins by 30th June 2012). Secondly, Food (C3) has the highest per venue tips, which indicates that Foursquare users are more likely to share their experiences of visiting food venues. Moreover, Residence (C7) and Professional & Other Places (C6) have the highest user stickiness, meaning that visitors to these venues tend to revisit the same venues often. On the other hand, the Food category has high per venue tips, but low user stickiness, so those venues attract users with fewer recurrent visits. This is easy to understand, since office, school, and church are three large subcategories of the category Professional & Other Places, and users checking in venues in these and Residence category tend to be people who live or routinely visit there, thus have higher stickiness. In comparison, venues in the Food and Travel & Transport categories tend to invite temporary visitors.

Summary: The above analysis implies that Travel & Transport (C9) is the most popular category in attracting per venue check-ins, whereas the Food (C3) category generates the most per venue tips. The Residence (C7) and Professional & Other Places (C6) categories attract users with the most recurrent visits.

C. Popularity analysis over venue ages

Foursquare was launched March 2009. By the time (June 30 2012) we finished our data collection, it was 40 months old. In this subsection, we study how venue ages affect their popularity, namely, whether old or new venues attract more people's interests. Figure 5 gives the numbers of new venues created from March 2009 to June 2012, in total and split for selected regions. Note that during Foursquare's first 14 months (Mar 2009 - Apr 2010), the total venue creation rate increased dramatically (the line has a steep slope), but after April 2010 venue creation stabilized to a steady rate (zero slope).

When examining the results by region, we observe that different regions have different patterns as shown in Figure 5 and the zoomed-in area for Dec 2009 to Jun 2012. To be precise, for Foursquare's initial set of regions (New York




Fig. 5. Number of new venues over time




Fig. 6. Average venue popularity over ages

City, Colorado, Los Angeles, etc), people have generated venues since day one. The number of new venues increased dramatically for the first 13 months, and then the rate stabilized for the next two years. For some regions that were added 7 to 11 months later (Singapore, Seoul, Pittsburgh, Wyoming, Orlando, Sydney, Paris, etc), their current venue creation rates are either stable or decreasing. On the other hand, in regions, such as Belgium, Switzerland, Slovenia, and Cairo, the rates of venue creation are still increasing.

We calculate the average number of check-ins per venue per month to capture the popularity of the venues with different ages (See Figure 6). First, we observe that older venues created at the early stage of Foursquare have higher popularity than newer venues. The venues created by Dec 2009 attract on average 50-120 check-ins per venue per month. The venues created during the later months (from Oct 2010 - Jan 2012)

generate on average only 3-4 check-ins per venue per month. However, the venues created in the latest 5 months have slightly higher popularity. We conjecture that this happens because people would like to try and visit new places, thus more check-ins and tips are generated during the first few months. In addition, we observe (Figure 6) that venues created in the 9th month (Nov 2009) attract far more check-ins, users, and tips than those venues created during the adjacent months. This happens because during Nov 2009, Foursquare twice launched new areas with 15 and 50 new cities on Nov 4 and Nov 19, respectively, which doubled their geographic coverage [5], [7]. This phenomenon implies that the first set of venues created in a newly launched area more likely have higher popularity. For example, Singapore and Orlando are two new areas launched on 19th Nov 2009, and the landmarks in these areas, such as VivoCity, Jurong Point, and SeaWorld, were created at that time, and they are very popular over time, namely, each of these venues attracts thousands of check-ins every month.

IV. RELATED WORK

LBSNs have attracted great interest to study large-scale location based datasets and their implications in improving location-based online services, e.g., user mobility prediction, friendship and venue recommendations, etc. [22], [23] study the correlation between the friendship relations and the user location relevance, and design a social link prediction system by incorporating the users' location information. [14], [21] incorporate both user preference and venue spatial relevance for venue recommendation in LBSNs. [15], [19], [24] examine large-scale (uncertain) trajectory data, with applications in revealing user mobility patterns and providing personalized route recommendations. [13] presents a measurement study of the temporal evolution of Gowalla, a location-based social network, and reveals interesting insights of how friendship and social triangles are established. Different from all these efforts, we study the key features that make venues popular, e.g., generating more visits (check-ins) or comments (tips) from users, and the characteristics popular venues usually possess. Our work is the first attempt to answer these questions.

V. CONCLUSION

In this paper, we investigate what drives the Foursquare venues to be popular, namely, attracting people to visit (check in) or leave comments (tips) to them. We exhaustively collected 2.4 million venues and their information from 14 regions around the world. The popularity of each venue is captured by the venue's statistics, such as the number of check-ins, users, and tips. By analyzing this unique dataset, we highlight three aspects that have significant impact on venue popularity, including the completeness of venue profile information, venue category, and venue age. First of all, a venue with more complete profile information is more likely to be popular, and "contact" and "cross street" are the two most influential attributes for venue popularity. Secondly, venues in the Food category attract the most (43%) public comments (tips) by users, and the Travel & Transport category is the most popular category with the highest per venue check-ins, i.e., each venue in this category attracts on average 376 check-ins. Moreover, venues in the Residence and Professional & Other Places categories have the most repeat visits (check-ins) from recurrent

users. Last but not the least, when looking into the venue ages, we discover that old venues in general are more popular than new venues. Our results shed light on understanding the factors affecting venue popularity, with applications in venue recommendations and targeted advertisement in location based social networks. These findings may help advertisers to select promising candidate venues for more effective advertisement placement, and venue owners to improve their venues' attraction to customers. We also made the anonymized version of our venue dataset available to the research community [10].

As part of our future work, we are interested in collecting, analyzing, and mining Foursquare dataset to study various user behaviors, such as community detection [18] and social influence propagation [16].

REFERENCES

- [1] A billion check-ins to foursquare venues. <http://blog.foursquare.com/2011/09/20/billion/>.
- [2] Facebook places. <http://touch.facebook.com/>.
- [3] Foursquare. <https://foursquare.com/>.
- [4] Foursquare API. <https://developer.foursquare.com/>.
- [5] Foursquare blog 2009/11. <http://blog.foursquare.com/2009/11/>.
- [6] Foursquare category tree. <https://developer.foursquare.com/docs/explore#req=venues/categories>.
- [7] Foursquare goes global with 50 new cities. <http://mashable.com/2009/11/19/foursquare-50-more-cities/>.
- [8] Foursquare house rules. <https://foursquare.com/info/houserules>.
- [9] Foursquare on wikipedia. <http://en.wikipedia.org/wiki/Foursquare>.
- [10] Foursquare venue dataset. <http://www.cs.umn.edu/~yanhua/dataset.html>.
- [11] Google Geo-coding API. <https://developers.google.com/maps/documentation/geocoding/>.
- [12] Google+ Local. www.google.com/+learnmore/local.
- [13] M. Allamanis, S. Scellato, and C. Mascolo. Evolution of a location-based online social network: Analysis and models. In *IMC'12*, 2012.
- [14] J. Bao, Y. Zheng, and M. F. Mokbel. Location-based and Preference-Aware Recommendation Using Sparse Geo-Social Networking Data. In *SIGSPATIAL GIS'12*, 2012.
- [15] E. Cho, S. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *KDD'11*, 2011.
- [16] Y. Li, W. Chen, Y. Wang, and Z.-L. Zhang. Influence diffusion dynamics and influence maximization in social networks with friend and foe relationships. In *WSDM '13*. ACM, 2013.
- [17] Y. Li, M. Steiner, L. Wang, Z.-L. Zhang, and J. Bao. Dissecting foursquare venue popularity via random region sampling. In *CoNEXT'12 student workshop*, 2012.
- [18] Y. Li, Z. Zhang, and J. Bao. Mutual or unrequited love: Identifying stable clusters in social networks with uni- and bi-directional links. *WAW '12*, 2012.
- [19] H. Liu, L.-Y. Wei, Y. Zheng, M. Schneider, and W.-C. Peng. Route discovery from mining uncertain trajectories. In *ICDM'11*, 2011.
- [20] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil. An empirical study of geographic user activity patterns in foursquare. *ICWSM'11*, 2011.
- [21] D. Saez-Trumper, D. Quercia, and J. Crowcroft. Ads and the city: Considering geographic distance goes a long way. In *ACM RecSys'12*, 2012.
- [22] S. Scellato, A. Noulas, R. Lambiotte, and C. Mascolo. Socio-spatial properties of online location-based social networks. *ICWSM'11*, 2011.
- [23] S. Scellato, A. Noulas, and C. Mascolo. Exploiting place features in link prediction on location-based social networks. In *KDD'11*, 2011.
- [24] L.-Y. Wei, Y. Zheng, and W.-C. Peng. Constructing popular routes from uncertain trajectories. In *KDD'12*, 2012.
- [25] G. K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, 1949.