

A picture of Instagram is worth more than a thousand words: Workload characterization and application

Thiago H. Silva^{*}, Pedro O. S. Vaz de Melo^{*}, Jussara M. Almeida^{*}, Juliana Salles[†], Antonio A. F. Loureiro^{*}

^{*}Computer Science, Universidade Federal de Minas Gerais, Belo Horizonte, MG, Brazil

[†]Microsoft Research, Redmond, WA, USA

{thiagohs, olmo, jussara, loureiro}@dcc.ufmg.br, jsalles@microsoft.com

Abstract—Participatory sensing systems (PSSs) have the potential to become fundamental tools to support the study, in large scale, of urban social behavior and city dynamics. To that end, this work characterizes the photo sharing system Instagram, considered one of the currently most popular PSS on the Internet. Based on a dataset of approximately 2.3 million shared photos, we characterize user’s behavior in the system showing that there are several advantages and opportunities for large scale sensing, such as a global coverage at low cost, but also challenges, such as a very unequal photo sharing frequency, both spatially and temporally. We also observe that the temporal photo sharing pattern is a good indicator about cultural behaviors, and also says a lot about certain classes of places. Moreover, we present an application to identify regions of interest in a city based on data obtained from Instagram, which illustrates the promising potential of PSSs for the study of city dynamics.

I. INTRODUCTION

Mark Weiser, in his classical article entitled “The computer for the 21st century” that appeared in the Scientific American magazine [1], popularized the concept of ubiquitous computing, which envisions the availability of a computing environment for anyone, anywhere, and at any time. It may involve many wirelessly interconnected devices, not just traditional computers, such as desktops or laptops, but may also include all sorts of objects and entities such as pens, mugs, phones, shoes, and many others. Although this is not the reality yet, and this concept has been extended to include, for example, the Internet of Things, much has been done in this direction in the past 20 years after the publication of Weiser’s seminal paper. In this scenario, Wireless Sensor Networks (WSNs) [2] play an important role since they are designed to collect data about the physical environment where they are inserted and provide such information to the end user or other entities. Moreover, there is an increasing use of participatory sensing systems (PSSs) [3], allowing people connected to the Internet to provide useful information about the context in which they are at any given moment.

Indeed, PSSs have the potential to complement WSNs in several respects. As WSNs are typically designed to sense areas of limited size, such as forests and volcanoes, PSSs can reach areas of varying size and scale, such as large cities, countries or even the entire planet [4]. Furthermore, a WSN is subject to failure, since its operation depends upon proper coordination of actions of its sensor nodes, which have severe hardware and software restrictions. On the other hand, PSSs are formed by independent and autonomous entities, i.e., humans, which make the task of sensing highly resilient to individual failures.

The success of PSSs is directly connected to the popularization of the *smartphone*, which became the most widely adopted personal computing device [5]. *Smartphones* have a rich set of built-in sensors, such as GPS, accelerometer, microphone, camera, gyroscope and digital compass, and typically remain turned on all time. However, sensing not only depends on the data generated by these sensors but also on the user’s subjective observations. Currently, there are several examples of PSSs already deployed and used by *smartphones*, such as Waze¹ to report traffic conditions in real time, and Weddar² to report weather conditions. Moreover, there are photo-sharing services, such as Instagram³, where users can send images in real time to the system. In particular, Instagram is currently one of the most popular PSS, with nearly 100 million users and more than 1 billion photos received, having every second a new user registered and 58 new photos uploaded [6].

The main objective of this work is to characterize the participation network of Instagram, aiming to show the challenges and opportunities emerging from participatory sensing performed by users of this application. Based on a dataset of approximately 2.3 million photos, we show the planetary scope of the network, as well as the highly unequal frequency of photos sharing, both spatially and temporally, which is highly correlated with the typical routine of people. Moreover, we also show how to design useful applications based on Instagram as we present a technique to identify regions of interest within a city. This application illustrates the immense potential of PSSs to study the dynamics of cities. To the best of our knowledge, this is the first study to characterize the use of Instagram by the photos shared.

The rest of this paper is organized as follows. Section II presents the related work. Section III discusses the participation of human beings in the process of sensing, addressing participatory sensing systems and participatory sensor networks (PSNs), arising from PSSs. Section IV presents the characterization of a PSN derived from Instagram. Section V describes an application to classify regions in a city using Instagram data. Finally, Section VI presents the conclusion and future work.

II. RELATED WORK

The process of sensing the environment may involve humans as (i) the target of the process [7], or (ii) the person responsible for collecting the data [8], [9]. In this paper, we focus

¹<http://www.waze.com>

²<http://www.weddar.com>

³<http://www.instagram.com>

on the second case, considering systems that employ mobile devices, such as *smartphones*, to build a participatory sensor network, which is described in Section III-B. In the literature, we can find several systems that consider the involvement of humans in the sensing process. Some of those participatory sensing systems (PSSs) include, for example, traffic [10] and noise [11] monitoring systems.

The success of PSSs depends mainly on the continuous participation of users along the time. Reddy et al. [12] propose incentive mechanisms based on micro-payments, which are small amounts of money given to the user when he/she performs certain activities in the system. Besides the continuous participation of a user, the system needs to ensure the quality of the shared data [13]. For example, in several PSSs users can fabricate false data supposedly sensed at a low cost. Therefore, data integrity is not always guaranteed [14].

There are several proposals devoted to the study of specific characteristics of PSSs. For example, in location sharing services like Foursquare, Cheng et al. [15] observe that users follow a pattern of mobility simple and feasible to be reproduced. In this direction, Cho et al. [16] observe that humans perform short trips that are periodic in space and time and are not affected by a social network structure, which, in its turn, influences only long distance trips.

Scellato et al. [17] show that 40% of social relations arising among users of three popular online location-based services happen within [100]km. Noulas et al. [18] analyze the dynamics of sharing in location sharing services and show, for example, that the distribution of the number of check-ins is highly uneven, being well modeled by a power law.

Other studies propose using data derived from PSSs in new applications, since this type of data helps to better understand the physical boundaries and notions of space [19]. In this direction, Cranshaw et al. [20] present a model to classify regions of a city based on patterns of collective activities, while Noulas et al. [21] propose using categories of places registered on Foursquare to classify areas and users of a city.

In a previous work [4], we analyze the properties of participatory sensor networks derived from two location sharing applications: Gowalla and Brightkite. We analyze the spatial and temporal distributions of check-ins performed by users of these systems to collect relevant evidence so we design new services and applications. In [22], we proposed a new way to visualize the dynamics of cities based on habits and routines of people collected from check-ins on Foursquare.

The only study about Instagram found is the one performed by Rainie et al. [23], where the authors interviewed Instagram users finding, for instance, that Instagram is more likely to be used by young adults. There are also studies that analyze similar photo sharing systems, such as Flickr, which is not a system accessed mostly by smartphones, i.e., it is not a conventional PSS. However, some of those studies take into account large scale geotagged photos, what make them particularly related to this work. Crandall et al. [24] study how to organize a large number of geotagged photos, combining analysis of tag text with geospatial data of the photo. Their goal was to estimate the location of a photo without considering the geospatial data. As result, their work reveals properties about landmarks of a city. Van Zwol [25] characterize the

interest of users in photos collected from Flickr. That work explores the spatial dimension to investigate the interest of users by a photo, showing that the geographic distribution is more focused around a geographic location.

Our work differs from previous ones (including ours) since it focus on a new system of great popularity nowadays – the Instagram. To the best of our knowledge, this is the first characterization of Instagram by photos shared by users. In fact, we specifically analyze the Instagram from a crowdsensing point of view. Moreover, continuing our recent studies [4], [22], this work examines the dynamics of cities across PSSs, showing that photo-sharing systems, particularly the Instagram, can also be used for mapping the characteristics of urban locations at a low cost.

III. HUMANS IN THE SENSING PROCESS

The focus of this paper is on systems that rely on humans' participation in the sensing process, where they are responsible for local data sharing. Such data can be obtained with the aid of sensing devices such as sensors embedded into *smartphones* (e.g., GPS) or by human sensors (e.g., vision), being subjective observations produced by them [8].

A. Participatory Sensing

Participatory sensing is the process where humans actively use mobile devices and cloud computing services to share local environmental data such as a picture [3]. It differs from opportunistic sensing [26] mainly by the user participation, which is minimal in the latter case. In this work, we consider that a fundamental point in participatory sensing is the user desire to share data, regardless of the process applied to generate it. In fact, we consider also manually user-generated observations. Participatory sensing with these characteristics is usually referred to as *ubiquitous crowdsourcing* [13] or *mobile crowdsensing* [27]. The popularity of participatory sensing systems has grown rapidly with the increasing use of cell phones embedded with sensors and the ubiquity of wireless access to the Internet access. These devices have become a powerful platform that includes capabilities of sensing, computing and communication.

A data sensed in a participatory sensing application is (i) obtained through physical sensors (e.g., GPS) or human observations (e.g., road congestion), (ii) defined in time and space, (iii) obtained automatically or manually, (iv) structured or unstructured, and (v) voluntarily shared or not. To illustrate this type of system, consider an application for traffic monitoring, such as Waze. Users can share comments about accidents or congestion manually. It is still possible to calculate the speed of the car and automatically share the car's route with the aid of the GPS. With speed measurements of different vehicles sampled in a particular area, it is possible to infer, for example, congestion. In this case, users manage the application, which was created for this purpose, and the sensed data is structured. But if users use a microblogging service, such as Twitter⁴, the sensed data is unstructured. For example, the user "John" sends a message "I am facing slow traffic near the entrance of the campus."

⁴<http://www.twitter.com>

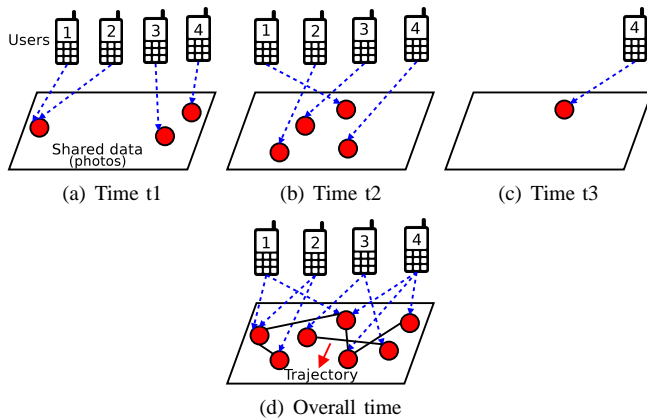


Fig. 1. Analyzed PSN: photo sharing service.

Photo-sharing services like Instagram are examples of participatory sensing applications. The sensed data is a picture of a specific place. We can extract information in many ways. For example, we can visualize in near real time how the situation is in a certain area of the city.

B. Participatory Sensor Network

In a Participatory Sensor Network (PSN), the user’s mobile device is the fundamental key element to obtain sensed data. Individuals carrying these devices are capable of sensing the environment and make relevant comments about it. Thus, each node in a PSN consists of a user with his/her mobile device. Similar to WSNs, the sensed data is transmitted to the server, or “sink node”. But unlike WSNs, PSNs have the following characteristics: (i) nodes are autonomous mobile entities, i.e., a person with a mobile device; (ii) the cost of the network is distributed among the nodes, providing a global scale; (iii) sensing depends on the willingness of people to participate in the sensing process; (iv) nodes transmit the sensed data directly to the sink; (v) nodes do not suffer from severe power limitations; and (vi) the sink node only receives data and does not have direct control over the nodes.


Figure 1 shows an example of a PSN comprised of photo-sharing services, which is analyzed in the following sections. Figures 1a, 1b, and 1c represent four users at different times. Photos shared by users are symbolized by a dotted arrow. Note that not all users perform activities at all times. After a certain interval, we can analyze the data in various ways. For example, Figure 1d shows a graph where the vertices represent the locations where the photos were shared and edges connect photos shared by the same user. With this graph, it is possible to obtain various results of interest, considering different parts of the world, providing a remarkable global infrastructure at a lower cost, as illustrated in Figure 2a.

IV. CHARACTERIZATION OF INSTAGRAM


In this section we analyze the participatory sensor network derived from Instagram.

A. Data Description

The data was collected via Twitter, which offers the possibility of integration with other platforms. This enables users to



(a) Number of photos n per pixel obtained from the value of ϕ shown in the figure, where $n = 2^\phi - 1$.



(b) Temporal variation of the number of photos shared by continent.

Fig. 2. Coverage of the PSN of Instagram.

announce photos available at Instagram, besides a plain text. In this case, photos of Instagram announced on Twitter become available publicly, which by default does not happen when the picture is published solely on the Instagram system.

Between June 30 and July 31 of 2012, we collected 2.272.556 tweets containing geotagged photos, posted by 482.629 users. Each tweet consists of GPS coordinates (latitude and longitude) and the time when the photo was shared.

B. Network Coverage

In this section, we analyze the coverage of the PSN of Instagram at different spatial granularities, starting around the planet, then by continents and cities and ending up at neighborhoods. Figure 2a shows the coverage on the planet by the PSN of Instagram as a heat map of user participation: darker colors⁵ represent larger numbers of photos shared in the particular area. Despite being a fairly comprehensive coverage on a planetary scale, it is not homogeneous. Figure 2b shows the number of photos shared by continent along the time. Note that the sensing activity in the Americas (North and South), Europe and Asia is at least an order of magnitude greater than in Africa and Oceania. Moreover, it can be observed that the participation of users in North America is slightly higher than in Latin America, Europe and Asia.

Now we evaluate the participation of users in Instagram in eight large and populous cities in five continents: New York/USA, Rio de Janeiro/Brazil, Belo Horizonte/Brazil, Rome/Italy, Paris/France, Sydney/Australia, Tokyo/Japan and

⁵Colors of the heat map for all subfigures are in the same scale.

Cairo/Egypt. Figure 3 shows the heat map of the sensing activity (photo sharing) in each one of these cities. Again, darker colors represent a greater number of pictures in a given area. We can observe a high coverage for some cities, as shown in Figures 3a (New York), 3e (Paris) and 3g (Tokyo). However, we can see in Figure 3h that the sensing in Cairo, which also has a large number of inhabitants, is significantly lower. Such difference in coverage may be explained by several factors. Besides the economic aspects, differences in the culture of the inhabitants of this city when compared with cultures present in the other cities analyzed may have a significant impact on the adoption and use of Instagram [28].

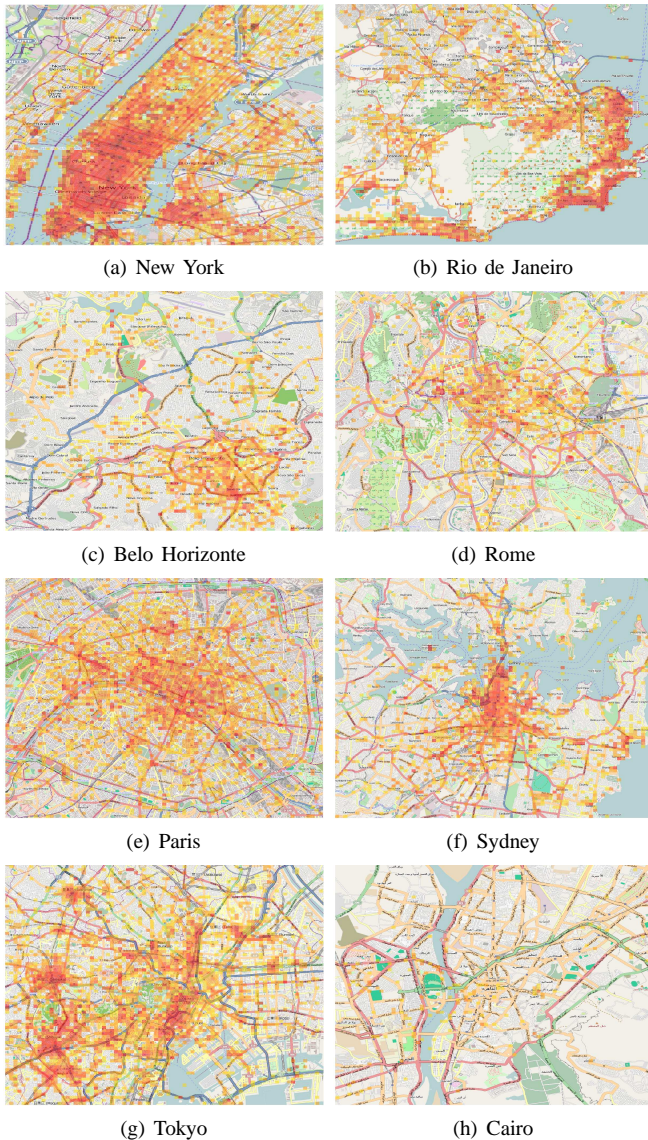


Fig. 3. Spatial coverage of Instagram in eight cities for all shared photos. The number of pictures in each area is represented by a heat map, where the scale varies from yellow to red (more intense activity).

Furthermore, we can see that the coverage in Rio de Janeiro and Sydney is more heterogeneous compared with the coverage in Paris, Tokyo and New York. This is probably because of the geographical aspects that these cities have in common, i.e., large green areas and large portions of water. Rio de Janeiro, for example, has the largest urban forest in the world,

located in the middle of the city, along with many hills that are not accessible to people. These aspects limit the geographic coverage of the sensing. Moreover, in both cities the points of public interest such as tourist spots and shopping centers are unevenly distributed throughout the city. There are large residential areas with few points of this type, while other areas have large concentrations of these points. These results are qualitatively similar to those reported in [4], [22] for PSNs derived from three location sharing applications and for different cities. This observation demonstrates the potential of Instagram as a tool for participatory sensing in large urban regions.




Fig. 4. Example of identification of a quadrant.

As the users' participation can be quite heterogeneous within a city, we propose to divide the area of the cities into smaller rectangular spaces, as in a grid. We call each rectangular area of a *quadrant* within a city and, from this, we analyze the number of photos shared in these quadrants. In this paper, we consider that a quadrant has the following delimitation: 10^{-4}° (latitude) \times 10^{-4}° (longitude). This represents an area of approximately 8×11 meters in New York City and 10×11 meters in Rio de Janeiro. For other cities, the areas can also vary slightly, but this does not affect the analysis. We believe that this is a reasonable size to represent an area of a venue, enabling then analysis of users' activity at venue level in a city. Figure 4 illustrates the process of dividing the area of a city in quadrants and how it is the association of geographic coordinate (24.0001433; 3.000253) to a quadrant X.

Figure 5 presents the complementary cumulative distribution function (CCDF) of the number of photos shared in a quadrant of the city of New York (Figure 5a) and all locations in our database (Figure 5b). First, note that in both cases, a power law describes well this distribution. This implies that most of the quadrants have few shared photos, while there are few areas with hundreds. These results are consistent with the results presented in [4], [18], which study the participation of users in location sharing systems. In systems for photo sharing, as well as systems for location sharing, it is natural that some areas have more activity than others. For example, in tourist areas the number of shared pictures tends to be higher than in a supermarket, although a supermarket is usually a location quite popular. If a particular application requires a more comprehensive coverage, it is necessary to encourage users to participate in places they normally would not. Micro-payments or scoring systems are examples of alternatives that might work in this case.

As previously shown, a PSN can have planetary scale coverage. However, it was also shown that such coverage can be quite heterogeneous, in which large areas are practically uncovered. Figure 6 shows the total network coverage considering the temporal dimension, i.e., the number of localities that are active (i.e., sensed) in a given time interval considering all available data. The maximum number of quadrants sensed




Fig. 5. Distribution of the number of photos in quadrants.

per hour corresponds to only approximately 0.2% of the total number of areas in our dataset (1.030.558). In other words, the instant coverage of the PSN of Instagram is very limited when we consider all locations that could be sensed on the planet. This means that the probability of a quadrant to be sensed on a random time is very low.




Fig. 6. Temporal variation in the number of sensed quadrants.

C. Sensing Interval

Participatory sensor networks are very scalable because their nodes are autonomous since users are responsible for their own operation and functioning. As the cost of infrastructure is distributed among the participants, this massive scalability and coverage is achieved more easily. The success of such a network is to have continuous participation with high quality. The sensing is efficient since users are kept motivated to share their resources and to sense data frequently.

Now we investigate the frequency in which users share photos in Instagram. Figure 7a shows the histogram of the inter-sharing time Δ_t between consecutive photos in a typical popular quadrant. Note that the histogram is well fitted by a log-logistic distribution [29] that has bursts of activity and long periods of inactivity: there are times when many photos are shared within a few minutes and there are times when there is no sharing for hours. This may indicate that the majority of photo sharing, in this popular area (as in others), occurs at specific intervals, probably related to the time when people usually visit them. For example, sharing photos in restaurants is likely to happen during lunch and dinner times. Applications based on this type of sensing should consider that the user participation can vary significantly along the time.

Another interesting observation related to the inter-sharing time Δ_t can be extracted from Figure 7b, which shows the

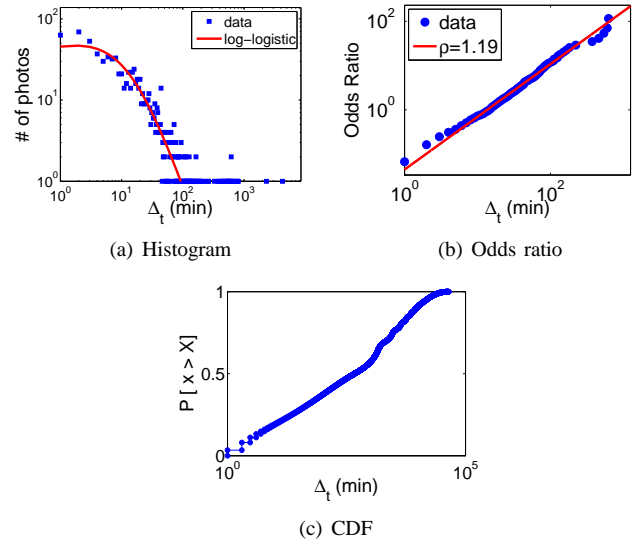


Fig. 7. Distribution of the time interval between shared photos in a popular quadrant.


odds ratio function (OR) of these intervals. The OR is a cumulative function where we clearly see the distribution behavior both in the head and in the tail. Its formula is given by $OR(x) = \frac{CDF(x)}{1-CDF(x)}$, where $CDF(x)$ is the cumulative density function, in this case, of the inter-sharing time Δ_t distribution. As in [30], the OR of the inter-sharing time between photos also presents a power law behavior with slope $\rho \approx 1$. This suggests that the mechanisms behind human activities can be simpler and more general than those proposed in the literature, because they depend on a lot of parameters [31]. Based on this fact and also on Figure 7c, we can observe that a significant portion of users performs consecutive photo sharing in a short time interval. About 20% of all observed sharing occurs within 10 minutes. As discussed in Section IV-E, this suggests that nodes tend to share more than one photo in the same area.

Related to this analysis, it is interesting to verify the feasibility of an application for near real-time visualization of a certain area of a city. For that, a central question is: what is the probability to obtain one picture of an area in a given time? To address this question, we select a popular area of our dataset (south of Manhattan), shown in Figure 8a, and divide it in eight sectors of equal size.

Figures 8b-e show the mean probability, along with its confidence interval of 95%, of seeing a picture in each of these sectors in the next 1-minute, 15-minutes, 30-minutes, and 60-minutes. All these probabilities are calculated for four different times of the day: dawn (Figure 8b), morning (Figure 8c), afternoon (Figure 8d), and night (Figure 8e). We observe that during the afternoon and night the difference between the probability of seeing a picture in the next 15 minutes and 60 minutes are not very high in most sectors. On the other hand, during the dawn and morning this difference is more expressive. This is explained by the low sharing frequency during the dawn and morning periods, as observed in Figure 9. Note also that even for a very popular area the probability to obtain a picture in the next minute is very low, for all four periods of the day. This means that applications that need a

considerable amount of photos within a small interval have to be aware that this may not be feasible.

The results in Figure 8 can also be used to better understand those sectors. For instance, Sector 8 seems to be the least popular among the others, despite the biggest part of water in that sector. If we analyze the probability of a photo in the next 15-minutes, we can also see that during the dawn, Sectors 3, 5, and 6 are the most popular ones, which might indicate that those sectors have a more intense nightlife. This information could be useful, for example, in a tourist guide, being one feature in an algorithm to recommend areas in a city.



(a) Sectors of NY

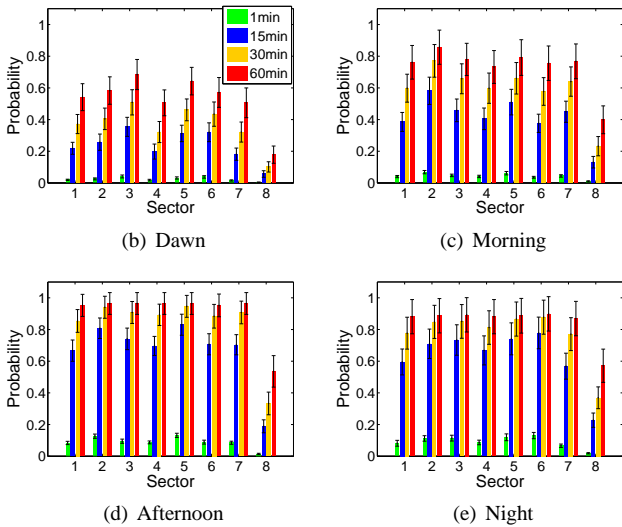


Fig. 8. Mean probability of obtain a picture in the next 1-minute, 15-minutes, 30-minutes, and 60-minutes, for eight popular areas during the dawn, morning, afternoon, and night.

D. Seasonality

We now analyze how humans’ routines affect the data sharing. First, we study all localities present in our dataset, and then we study the sharing pattern for some cities from different continents.

1) *All Localities*: Figure 9a shows the weekly pattern of photo sharing in Instagram⁶. As expected, the network participation presents a diurnal pattern, implying that the overnight sensing activity is quite low.

⁶The time of sharing was normalized according to the location where the photo was taken.

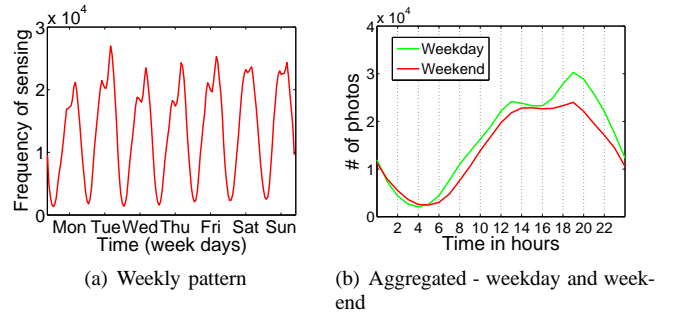


Fig. 9. Temporal photo sharing pattern.

Considering weekdays, we can see a slight increase in activity throughout the week, except for Tuesday, when there is a peak of activity. Cheng et al. [15] analyze location sharing systems and observe the same behavior. This suggests that during the period of data collection, an unusual event may have happened on Tuesday that resulted in an abnormal number of shared photos. Finally, observe two peaks of activity throughout the day, one around lunch and the other at dinner time. Unlike the behavior observed for location sharing [15], for photo sharing there is no peak of activity at breakfast time.

We also analyzed the behavioral patterns during weekdays and weekends. Figure 9b shows the average number of photos shared per hour during weekdays (Monday to Friday), and also during the weekend (Saturday and Sunday). As we can see, the peaks during weekdays happen around 13:00 (lunch) and 19:00 (dinner), but on weekends there is no peak of activity at lunchtime. Rather, the activity remains intense throughout the afternoon until early evening, with a slight increase at 19:00.

2) *Selected Areas*: We now turn our attention to the photo sharing pattern throughout the day in Rio, Sao Paulo, Osaka, Tokyo, Barcelona, Madrid, Chicago and NY during weekdays and weekends. These results are shown in Figure 10⁷. It is interesting to note that, even when we analyze separate cities, we still do not observe, for most of the cities, a clear peak of photo sharing around the breakfast time, as observed for location sharing [15].

Studying weekdays first, we can see that cities from Japan (Figure 10c), Spain (Figure 10e) and USA (Figure 10g) present peaks of photo sharing that reflect typical lunch and dinner times. On the other hand, not all peaks in the Brazilian curves (Figure 10a) represent typical meal times. This might indicate that Brazilians share photos in uncommon moments. We conjecture that the peak of 6:00pm is due a “happy hour” and the peak of 9:00pm is due to a leisure activity that happens in a pub, theater, concert, etc. Another difference is that, in general, the Brazilian activity is more intense late at night. During weekdays it is possible to observe a certain similarity of sharing patterns between Japanese, Spanish, and American cities.

However, during the weekends these patterns are very distinct. The Brazilian curve still presents an unusual peak at 5:00pm and the Spanish and American curves now present more intense activity around the “brunch”/lunch time. These

⁷Each curve is normalized by the maximum number of photos shared in a specific region representing the city.

observed patterns might express cultural behaviors of inhabitants of those countries, presenting somehow the signature of a certain culture. This hypothesis is reinforced because we surprisingly see that the pattern for each city in the same country is fairly similar on weekdays, and also on weekends, at the same time, being distinct from patterns observed for other countries.

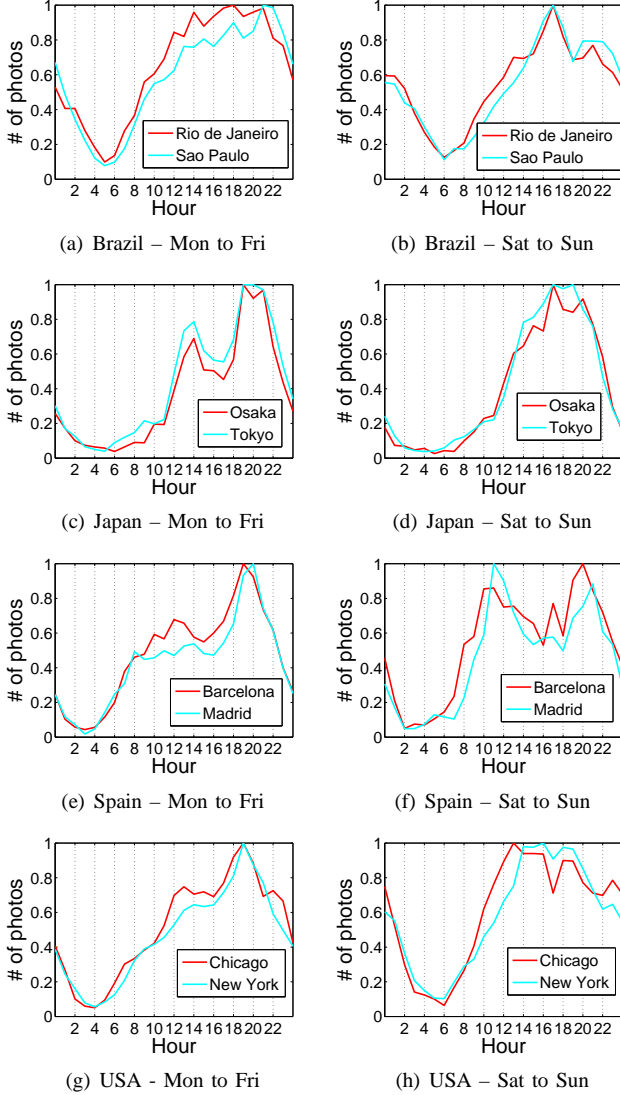


Fig. 10. Photo sharing throughout the day in Rio, Sao Paulo, Osaka, Tokyo, Barcelona, Madrid, Chicago and NY.

E. Node Behavior

In this section we analyze the sensing activity of each individual node (i.e., user plus *smartphone*) in the PSN. Figure 11 shows that the distribution of the number of photos shared by each user of our database has a heavy tail, meaning that user participation may vary widely. For example, about 40% of users contribute with only a photo during the considered period, while only 17% and 0.1% of users contribute more than 10 and 100 photos, respectively.

We also analyze the geographical distance between two consecutive photos shared by the same user, according to the

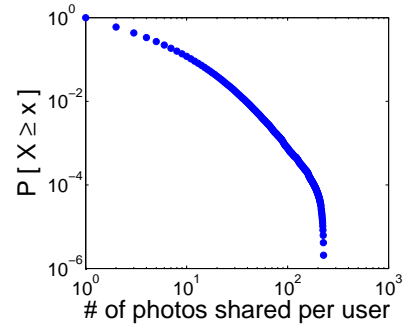


Fig. 11. Distribution of the number of photos shared by people.




Fig. 12. Distribution of the geographical distance between consecutive pictures of the same person.

geographic coordinates associated with each photo. Figure 12a shows the cumulative density function of the geographic distance between each pair of consecutive photos shared by each user in our dataset. It can be observed that a significant portion (about 30%) of the distances between consecutive photos are very short (less than [1]meter). This indicates that users tend to share multiple photos in the same location. This hypothesis is reinforced by the significant portion of time intervals between consecutive pictures of short duration shown in Figure 7c: 20% of these intervals (Δ_t) do not exceed 10 minutes. This was not observed in the same proportion for location sharing. Noulas et al. [18] observe that 20% of the shared locations happen up to [1]km away. For shared photos, this value is approximately 45%. This result can be explained by the simple fact that a photo can contain much more information than one location. For example, in a restaurant users could share photos of his/her friends at the place, food, or a particular situation, but tend to share their location only once.

We now analyze each user separately. Figure 12b shows the distribution of the median distance between consecutive sharing computed for each user. Note that at least 50% of consecutive photos of a significant portion of users (about 20%) are taken at a very short distance (around [1]meter).

Finally, we study the performance of nodes considering the total traveled distance, the coverage in the city of NY, and total number of contributed photos. To analyze the coverage, we consider the area of NY (Figure 13a), which was divided into 27 sectors of equal size. Figure 13b shows a 3-D plot for the three dimensions considered. We are able to observe the existence of “super nodes” in the system, indicated by a green circle. These nodes share a lot of photos, travel long distances, and visit many different areas in the city (observed by the number of unique visited sectors). The identification

of this type of users is important for several reasons. As the success of a PSN relies on a continuous contribution, it is interesting to award this type of user to keep them active in the network. Besides that, nodes of this type might be good candidates to be selected, for example, in a network for information dissemination a city.

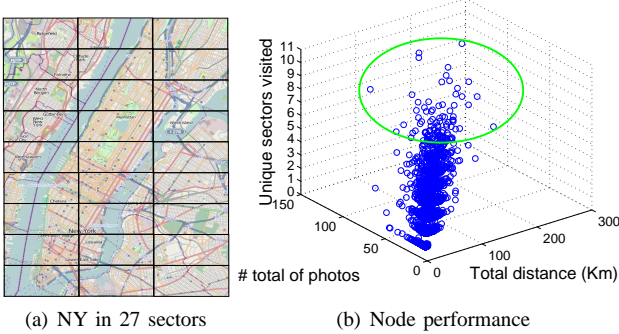


Fig. 13. Contribution of nodes, distance traveled, and coverage.

V. APPLICATION

It is quite common to find particular areas in a city that attract more attention of residents and visitors, here called *points of interest* (POI). Among the most visited POIs, we can mention the sights of the city. However, not all POIs are sights of a city. For example, an area of bars can be quite popular among city residents, but not among tourists. Furthermore, POIs are dynamic, in other words, areas that are popular today may not be tomorrow.

An application that naturally emerges from analyzing Instagram data is related to the identification of POIs in a city. This is possible because each picture represents, implicitly, an interest of an individual at a given moment. So, when many users share photos in a particular location at a given moment, it can be inferred that this place is a POI (note Figure 5).

More specifically, we formalize the process of identifying POIs by the following steps:

- 1) Each pair i of coordinates (longitude, latitude) $(x, y)_i$ is associated with a point p_i ;
- 2) calculate the distance [32] between each pair of points (p_i, p_j) ;
- 3) group all the points p_i that have a distance smaller than [250]m into a cluster C_k . This distance threshold was obtained by the method Complete-Linkage [33]. The result of this procedure is shown in Figure 14a, in which different colors represent different clusters k for the city of Belo Horizonte;
- 4) for each cluster C_k , we consider only one point (photo) per user. With that, the popularity of a cluster is now based on the number of different users that shared a photo in the cluster area. This procedure avoids considering areas visited by very few users, e.g., homes, as popular ones;
- 5) finally, for each cluster C_k , we create an alternative cluster C_r . Then, for each photo f_i , we randomly choose an alternate cluster C_r and we assign f_i to C_r . The number of photos assigned to each cluster

from that process follows a normal distribution with mean μ and standard deviation σ . Thus, from the original clusters C_k found in the previous step, we exclude those in which the number of photos is within the distance 2σ from the average μ , or is in the range $[\mu - 2\sigma; \mu + 2\sigma]$. The idea of this step is to exclude those clusters that may have been generated by random situations, i.e., those that do not reflect the dynamics of the city.

Figure 14b shows POIs obtained through this process. Observe the significant smaller number of points compared with the ones shown in Figure 14a. Besides identifying POIs in a city, we can also separate the sights from POIs. For this, first we generate a graph $G(V, E)$, where the vertices $v_i \in V$ are all POIs and there is an edge (i, j) from the vertex v_i to the vertex v_j if in a given time a user shared a photo on a POI v_j , after having shared a photo on POI v_i .

The weight $w(i, j)$ of an edge represents the total number of transitions performed from POI v_i to POI v_j considering transitions of all users. To identify sights, we consider that most tourists follow a well-known path within the city, being guided by the main sights of it. Moreover, at each point of interest he/she takes one or more photos and goes to the next tourist spot. Thus, we consider that edges (i, j) with high weights $w(i, j)$ denote these frequent transitions from one sight to another in a city.

After this, we exclude from G all edges (i, j) with weights $w(i, j)$ smaller than a threshold t , which is given by the probability of generating $w(i, j)$ randomly in a random graph $G_R(V, E_R)$. The identification of the value that separates edges with high weights from low weights is made as follows. First, we create a random graph $G_R(V, E_R)$ containing the same nodes of G . Then, for each sequence of n_u photos $f_u^1, f_u^2, \dots, f_u^{n_u}$ of each user u , we randomly assign a POI to each photo, what generates the random edges E_R of G_R . Thus, the sequence of locations where the photos were taken is random, but the total number of photos that were taken is preserved. The idea is to simulate random walks in a city. In this random fashion, the distribution of edge weights follows a normal distribution $N_w(\mu_w, \sigma_w)$ with mean μ_w and standard deviation σ_w .

When the probability p_w of generating an edge weight $\geq w_t$ in $G_R(V, E_R)$ is, according to $N_w(\mu_w, \sigma_w)$, close to zero, then all transitions $v_i \rightarrow v_j$ with $w(i, j) \geq w_t$ are popular, in which, according to our conjecture, are transitions between sights. For our dataset, the value of w_t which provides a probability p_w close to 0 is $w_t = 10$. As we can see in Figure 14c, the vertices (POIs) of the resulting graph represent practically all the sights of Belo Horizonte. The areas of the resulting POIs cover seven out of all the eight Landmarks recommended by TripAdvisor⁸ as the most important cultural and leisure areas of Belo Horizonte.

Notice the difference between Figures 14b and 14c, the first containing all POIs and the second only the sights of the city of Belo Horizonte. This means that inhabitants could also use this application to explore the city. Again, this application is interesting because it is able to identify POIs in

⁸www.tripadvisor.com

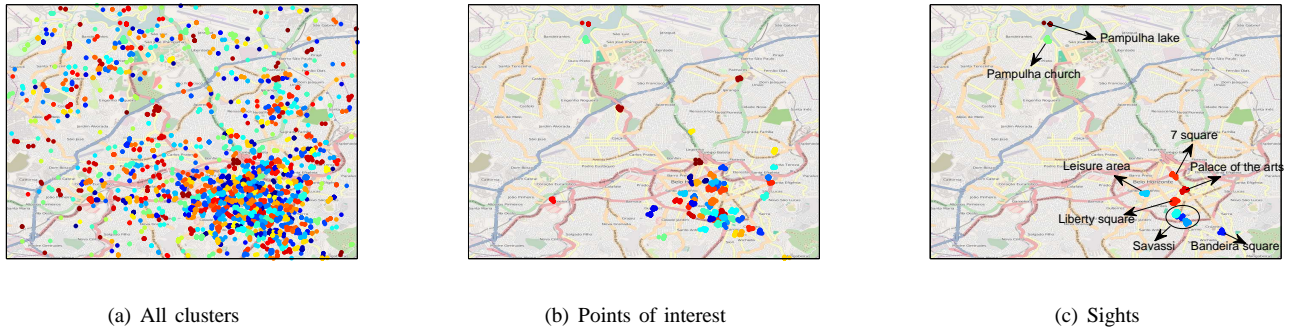


Fig. 14. Points of interest of Belo Horizonte.

a spatio-temporal context, which is fundamental, since POIs are dynamic and change over time.

A. The Vibe of POIs

Figures 14b and 14c show that a particular area (southeast) of the city has a high concentration of POIs. This can be useful to guide tourists in the city, for example, when choosing a hotel location. Another interesting information for city explorers is the time when certain POI is more popular. Intuitively, we know that certain types of places are frequented by people only at specific times of particular days. Figure 15 shows the number of shared photos per hour for all days of our dataset in different types of places. Figure 15a shows a soccer stadium. In that figure, the word WD indicates that the delimitation for dashed lines represents a weekend, five in total. All the activities shown represent games that happened during the analyzed interval. Observe also the lack of activity between games, indicating that this is an *event-oriented* POI. Other types of POIs are also event-oriented: night clubs (Figures ref:fig:vivePOIsb and 15c), and a convention center (Figure 15d). Note that the activities in night clubs concentrate more during weekends, on the other hand in a convention center most of the activity happens during weekdays.

Concerning other types of POIs, we can see in Figures 15e and 15f that people share photos in a mall in many different times of the day, during weekdays and weekends. This is expected due to the high number of different attractions that a mall usually offers every day of the week. We also show the frequency of two of the most famous touristic attractions of Belo Horizonte in Figures 15g and 15h. The sharing pattern in touristic spots are not as intense as POIs with a high concentration of people and attractions such as malls, or as periodic as an event oriented POI, such as night clubs. These are powerful features for classifying POIs by their type and suggesting users about the best time and day to make a visit to it.

Finally, as we can see, the temporal photo sharing pattern presents somehow a signature of POIs, meaning that may be possible to automatically identify anomalous events. This can be used to capture in near real time unexpected events, such as an accident, or an event happening in an unusual place, for instance a street party or a concert on a park. After identifying those events, we could use the shared pictures to check, in near real time, snapshots of those events. Figure 15a illustrates the potential of this application, showing some pictures for the

greatest peak of activity in that POI. In this case, a user could be aware that this event is a game of the Cruzeiro soccer team.

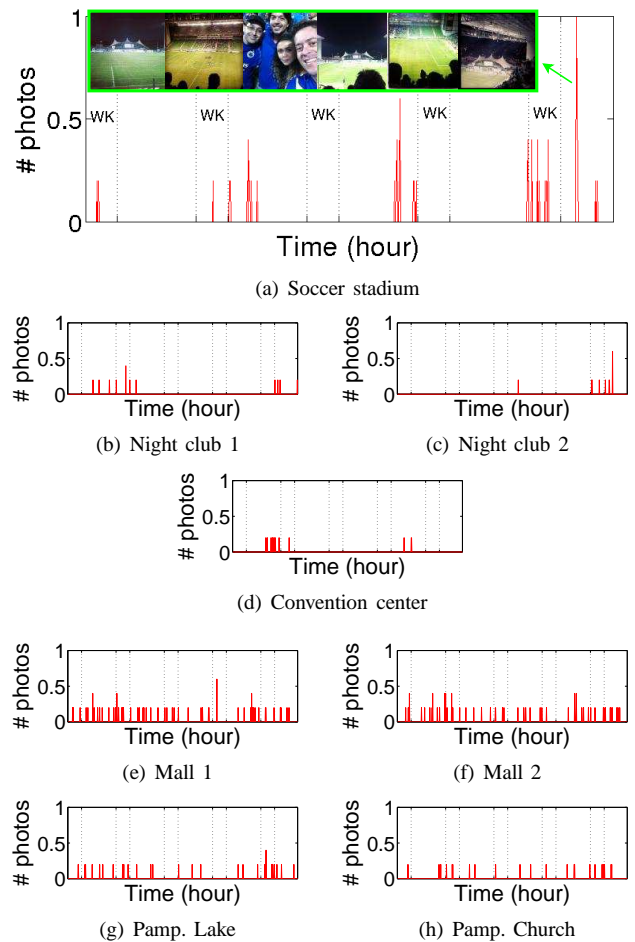


Fig. 15. The temporal photo sharing pattern for different types of POIs.

VI. CONCLUSION AND FUTURE WORK

In this work, we presented to the best of our knowledge, the first characterization of Instagram analyzing photos shared by the users. We analyzed the system treating it as a participatory sensing system. Thus, we discuss the spatial and temporal coverage of this network showing its global coverage. We observed that the frequency of photo sharing is spatio-temporal,

very unequal and correlated with routine human activities. We also observed that the temporal photo sharing pattern is surprisingly a good indicator about cultural behaviors, and also says a lot about certain classes of places. We also discuss an application that demonstrates the potential of a PSN derived from Instagram for studying the dynamics of cities.

As future work, we intend to analyze other PSNs and develop new applications that exploit these networks. For example, we imagine applications that jointly consider data from other participatory sensing systems, such as Waze (traffic conditions) and Weddar (weather), also considering different categories/interests of people.

REFERENCES

- [1] M. Weiser, "The Computer in the 21st Century," *Scientific American*, vol. 265, no. 3, pp. 94–104, 1991.
- [2] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, 2002.
- [3] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *Workshop on World-Sensor-Web (WSW'06)*, Boulder, Colorado, USA, 2006, pp. 117–134.
- [4] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, and A. A. F. Loureiro, "Uncovering Properties in Participatory Sensor Networks," in *Proc. of the 4rd ACM Int. Work. on Hot Top. in Planet-scale Meas. (HotPlanet'12)*, Lake District, UK, June 2012.
- [5] J. Krumm, *Ubiquitous Computing Fundamentals*. Chapman & Hall/CRC, 1st ed., 2009.
- [6] K. Daniells, "Infographic: Instagram statistics 2012," *Digital Buzz Blog*, 2012.
- [7] E. C. Larson, T. Lee, S. Liu, M. Rosenfeld, and S. N. Patel, "Accurate and privacy preserving cough sensing using a low-cost microphone," in *Proceedings of the 13th international conference on Ubiquitous computing*, ser. UbiComp'11. Beijing, China: ACM, 2011, pp. 375–384.
- [8] M. Srivastava, T. Abdelzaher, and B. Szymanski, "Human-centric sensing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 176–197, Jan. 2012. [Online]. Available: <http://dx.doi.org/10.1098/rsta.2011.0244>
- [9] M. Goodchild, "Citizens as sensors: The world of volunteered geography," *GeoJournal*, vol. 69, no. 4, pp. 211–221, 2007.
- [10] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell, "Bikenet: A mobile sensing system for cyclist experience mapping," *ACM Transactions on Sensor Networks*, vol. 6, no. 1, 2010.
- [11] R. K. Rana, C. T. Chou, S. S. Kanhere, N. Bulusu, and W. Hu, "Ear-phone: an end-to-end participatory urban noise mapping system," in *Proc. of the 9th ACM/IEEE Int. Conf. on Infor. Proc. in Sensor Networks*, ser. IPSN '10. Stockholm, Sweden: ACM, 2010, pp. 105–116.
- [12] S. Reddy, D. Estrin, M. Hansen, and M. Srivastava, "Examining micropayments for participatory sensing data collections," in *Proc. of the 12th ACM int. conf. on Ubiquitous computing (UbiComp '10)*. Copenhagen, Denmark: ACM, 2010, pp. 33–36.
- [13] A. J. Mashhadi and L. Capra, "Quality Control for Real-time Ubiquitous Crowdsourcing," in *Proc. of the 2nd Int. Workshop on Ubiquitous Crowdsourcing (UbiCrowd'11)*, Beijing, China, 2011, pp. 5–8.
- [14] S. Saroiu and A. Wolman, "I am a sensor, and i approve this message," in *Proc. of the 11th Work. on Mobile Comp. Systems and App.*, ser. HotMobile '10. Annapolis, Maryland: ACM, 2010, pp. 37–42.
- [15] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui, "Exploring Millions of Footprints in Location Sharing Services," in *Proc. of the Fifth Int'l Conf. on Weblogs and Social Media (ICWSM'11)*, Barcelona, Spain, 2011.
- [16] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: user movement in location-based social networks," in *Proc. of the 17th ACM Int. Conf. on Know. Disc. and Data Min. (KDD'11)*, San Diego, California, USA, 2011, pp. 1082–1090.
- [17] S. Scellato, A. Noulas, R. Lambiotte, and C. Mascolo, "Socio-spatial Properties of Online Location-based Social Networks," in *Proc. 5th International Conference on Weblogs and Social Media (ICWSN'11)*, Barcelona, Spain, 2011.
- [18] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil, "An Empirical Study of Geographic User Activity Patterns in Foursquare," in *Proc. of the Fifth Int'l Conf. on Weblogs and Social Media (ICWSM'11)*, Barcelona, Spain, 2011.
- [19] M. Bilandzic and M. Foth, "A review of locative media, mobile and embodied spatial interaction," *International Journal of Human-Computer Studies*, vol. 70, no. 1, pp. 66–71, Jan. 2012.
- [20] J. Cranshaw, R. Schwartz, J. I. Hong, and N. Sadeh, "The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City," in *Proc. 6th International Conference on Weblogs and Social Media (ICWSN'11)*, Barcelona, Spain, 2012.
- [21] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil, "Exploiting Semantic Annotations for Clustering Geographic Areas and Users in Location-based Social Networks," in *Proc. 5th International Conference on Weblogs and Social Media (ICWSM'11)*, Barcelona, Spain, 2011.
- [22] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, and A. A. F. Loureiro, "Visualizing the invisible image of cities," in *Proc. of IEEE Int. Conf. on Cyber, Phy. and Social Comp. (CPScom'12)*, Besancon, France, Nov. 2012.
- [23] L. Rainie, J. Brenner, and K. Purcell, "Photos and Videos as Social Currency Online," Pew Research, Tech. Rep., Sep 2012.
- [24] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg, "Mapping the world's photos," in *Proceedings of the 18th international conference on World wide web*, ser. WWW '09. Madrid, Spain: ACM, 2009, pp. 761–770.
- [25] R. van Zwol, "Flickr: Who is looking?" in *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, ser. WI '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 184–190.
- [26] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *Comm. Mag.*, vol. 48, no. 9, pp. 140–150, Sep. 2010.
- [27] R. Ganti, F. Ye, and H. Lei, "Mobile crowdsensing: current state and future challenges," *Communications Magazine, IEEE*, vol. 49, no. 11, pp. 32–39, november 2011.
- [28] F. Barth, *Ethnic groups and boundaries: the social organization of culture difference*, ser. Scandinavian university books. Little, Brown, 1969.
- [29] P. R. Fisk, "The graduation of income distributions," *Econometrica*, vol. 29, no. 2, pp. 171–185, 1961.
- [30] P. O. S. Vaz de Melo, C. Faloutsos, and A. A. Loureiro, "Human dynamics in large communication networks," in *Proc. SDM*, Mesa, Arizona, USA, 2011.
- [31] R. D. Malmgren, D. B. Stouffer, A. E. Motter, and L. A. N. Amaral, "A poissonian explanation for heavy tails in e-mail communication," *Proc. National Academy of Sciences*, vol. 105, no. 47, pp. 18 153–18 158, November 2008.
- [32] R. W. Sinnott, "Virtues of the Haversine," *Sky and Telescope*, vol. 68, no. 2, pp. 159+, 1984.
- [33] T. Sørensen, "A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons," *Biologiske Skrifter*, vol. 5, no. 4, 1948.