Large Scale Study of City Dynamics and Urban Social Behavior Using Participatory Sensing

Thiago H. Silva, Pedro O. S. Vaz de Melo, Jussara M. Almeida, Antonio A. F. Loureiro Computer Science Department, Universidade Federal de Minas Gerais, Belo Horizonte, MG, Brazil
{thiagohs, olmo, jussara, loureiro}@dcc.ufmg.br

***Abstract*—The ubiquitous availability of computing technology such as smartphones, tablets and other easily portable devices, and the worldwide adoption of social networking sites make it increasingly possible for one to be connected and continuously contribute to this massively distributed information publishing process. In this scenario, people act as social sensors, voluntarily providing data that capture their daily life experiences, and offering diverse observations on both the physical world (e.g., location) and the online world (e.g., events). This large amount of social data can provide new forms of valuable information that are currently not available, at this scale, by any traditional data collection methods, and can be used to enhance decision making processes. In this article, we argue that location-based social media systems, such as Instagram and Foursquare, can act as valuable sources of large scale sensing, providing access to important characteristics of urban social behavior much more quickly than traditional methods. We also discuss different applications and techniques that can exploit the data shared in these systems to enable large scale and near real time analyses and visualization of different aspects of city dynamics.**

***Keywords*-Urban computing; participatory sensing; location-based social media; Foursquare; Instagram; mobile social networks.**

I. Introduction

Mobile phones play a fundamental role in today's technologically-advanced community allowing people to communicate (almost) anywhere in the world and share all kinds of contextual information (e.g., location and opinion). Modern mobile phones, namely smartphones, are the new frontier for accessing the Internet and the World Wide Web. They are being manufactured with an increasing number of powerful embedded sensors of different categories (e.g., acoustic, sound, and magnetic vibration), enabling a variety of new applications and services. Indeed, smartphones are being used for many personal sensing applications, such as for monitoring physical exercises, and for wide participatory sensing applications, which are not limited to a particular individual (e.g., traffic conditions and noise pollution) [2].

Participatory sensing aims at monitoring large scale phenomena and require the active involvement of people to voluntarily share contextual information and/or make their sensed data available. Participatory sensing finds an effective platform for large scale reach in the increasing popularity of location-based social media applications, such as Instagram[1] and Foursquare[2], which combine the features of online social networks with location-based services.

These applications enables the observations of the actions of hundreds millions of people in large scale urban areas in (near) real time and over extended periods of time, helping us to better understand city dynamics and urban behavior worldwide, as well as to more quickly react to changes. This opens an unprecedented opportunity to revolutionize the way social science is done. Unlike traditional methods that rely on survey data, new techniques can be designed to exploit participatory data, which is much cheaper, more dynamic as it reflects current situations in (near) real time, and, more important, can easily reach planetary scale. Moreover, as we argue here, such participatory sensing applications have the potential to be a fundamental tool to better understand human urban interaction in the future, leveraging our awareness to different aspects of our lives in urban scenarios.

The goal of this article is to discuss the potential of location-based social media systems as sources of large scale participatory sensing from which valuable knowledge about city dynamics and urban social behavior can be drawn. To illustrate such potential, we present different methods and techniques that exploit the data shared in these applications to enable large scale and near real time analyses of different aspects of user behavior as well as discuss the technical challenges involved in building and deploying such methods.

II. Participatory sensor networks

Location-based social media (LBSM) systems, such as Foursquare and Instagram, build new virtual environments that integrate user interactions. Recently, due to the widespread adoption of smartphones, such systems are becoming increasingly popular, offering unprecedented opportunities of access to planetary scale sensing data. In traditional wireless sensor networks, the high costs associated with building and managing large scale topologies are prohibitive. In contrast, LBSM systems allow people to share useful data about the context they are inserted in at any time, making them potential sources of sensing at global scale. Such systems have been called participatory sensing systems (PSSs) [1], [2]. Participatory sensor networks (PSNs) can be derived from such systems, where each node represents a user able to sense context data with his/her mobile device. For example, in a PSN derived

[1]http://www.instagram.com.

[2]http://www.foursquare.com.

from Instagram, the sensed context data is a picture of a specific place where the user is located.

Figure 1 illustrates a PSN built from users with their portable devices sending sensed data about their locations to PSSs. The figure shows the sharing activities (represented by red dots) of four users at three different points in time (labeled “Time 1”, “Time 2”, and “Time 3”). Note that a user does not necessarily participate in the system at all times. After a given time, we can analyze this data in different ways. For instance, the bottom rightmost portion of the figure shows, as an aggregated view, a directed graph with nodes representing locations where data was shared and edges connecting locations that were shared by the same user. Using this graph we can extract, for instance, user mobility patterns. In fact, knowledge discovery in PSNs walks together with a wide range of studies that use graph theory for social network analysis (SNA) [3]. As we show in Section III, well known techniques used for SNA may be directly applied to analyze social oriented graphs derived from PSNs.

Thus, PSNs are an example of the interplay between technological networks and social networks, since a key element in a PSN is the human being. Many questions arise from this emerging concept. What are the properties of PSNs? What types of applications can we apply PSNs in? What are their limitations? Since PSNs reflect the dynamics of humans in terms of their routines, preferences and demographics, PSNs offer a new and cheap platform for studying human behavior in near real time and planetary scale. As the data provided by PSNs may be very complex, a fundamental step in any investigation is to characterize the collected data in order to understand its challenges and usefulness.

In this direction, we have studied the properties of PSNs derived from two Instagram datasets – Instagram-1 and Instagram-2, two from Foursquare – Foursquare-1[3] and Foursquare-2, one from Brightkite, and one from Gowalla, comprising over 30 million registers (photos or check ins) [1], [2]. Data, check-in or photo, have the following attributes: [user id; time; latitude; longitude]. The dataset Foursquare-1 has 4,672,841 check-ins, collected in 4/2012 (1 week), and this dataset in particular has extra attributes: [venue id; main category of the place; subcategory of the place]. The dataset Foursquare-2 has 4,548,941 check-ins, collected from 05/11/13 to 05/25/13. The dataset Instagram-1 has 2,272,556 photos, collected from 06/30/12 to 07/31/12, and the dataset Instagram-2 has 1,855,235 photos, collected from 05/11/13 to 05/25/13. Note that the datasets Foursquare-2 and Instagram-2 were collected at the same time.

Figure 2 summarizes some of the properties extracted from the analyzed PSNs. These results are representative of all systems. Figure 2a shows the coverage of a PSN derived from the Foursquare-2 dataset, where the number of locations n per pixel is given by the value of ϕ displayed in the colormap, where $n = 2^{\phi} - 1$. Clearly, the coverage is very comprehensive in a planetary scale. Figure 2b presents the complementary cumulative distribution function (CCDF) of the number of times someone shared data in a given location in Instagram-1 and Foursquare-1 datasets. For Instagram and Foursquare-2, a location is defined as an area of approximately 10x10 meters, whereas for the Foursquare-1, it is a registered venue. Consistently with previous observations [4], a power law distribution[4] fits well the data in both datasets, as most locations have only a handful of check-ins while a few locations have hundreds or even thousands of them. The intrinsic characteristics of different locations play an important role in determining the frequency of sharing in these locations. For example, users are more likely to share data in a restaurant than in a supermarket, regardless of their popularities. This imposes challenges if one needs a comprehensive and uniform data contribution.

Figure 2c shows the percentage of distinct locations where users shared data in a given time interval for the Instagram-2 and Foursquare-2 datasets, which have 598,397 and 725,419 locations respectively. The maximum percentage of distinct locations shared per hour is below 3% for all systems. This indicates that the short-term coverage of these PSNs is very limited considering all reachable locations, i.e., the probability of a random location being active in a given day is small.

Finally, Figure 2d illustrates how the seasonal behavior of humans affects data sharing in the Foursquare-1 dataset[5]. It shows the average number of check-ins on each hour of the day, for weekdays and weekends. As expected, user participation presents a diurnal pattern. Moreover, three peaks reflecting typical meal times - breakfast, lunch, and dinner - can be observed on weekdays, but not on weekends, consistently with typical routines of people.

These results illustrate the potential of PSN analysis to foster the study of urban social behavior. Although these results had been separately shown in previous work [1], [2], [5], we here argue that the observed properties are common to different PSNs - the same properties are observed on data collected from different applications, as well from the same application in different time intervals. More broadly, these properties reveal the potential of PSNs to drive various studies on city dynamics and urban social behavior, as shown in [4], [6]–[9], which have also exploited such properties to build services in different areas, including tourism and smart cities. Next, building upon previous observations, we present several techniques that make use of PSN analysis to drive large scale studies of city dynamics and urban social behavior.

III. Understanding City Dynamics and Urban Social Behavior

What is the current best way to study the dynamics of a city? How can we learn about the routines of its citizens, their movement patterns, its points of interest, and its cultural and economical aspects? One might choose to rely on official census data, while others may opt to simply get one or more guide books at their favorite bookstore. Although we are very fond to books and census efforts, do they always offer accurate

[3]Collected using one average machine (AMD FX 3300MHz with 16Gb RAM), the processing time of the average number of check-ins per day in New York City, 10K check-ins, took around 10 seconds.

[4]A power law distribution is characterized by a scaling parameter α, which is the slope of the curve in logarithm scale on both axes. The values of α for both datasets are shown in Figure 2b.

[5]Timestamps were normalized according to the local time of the check-in.

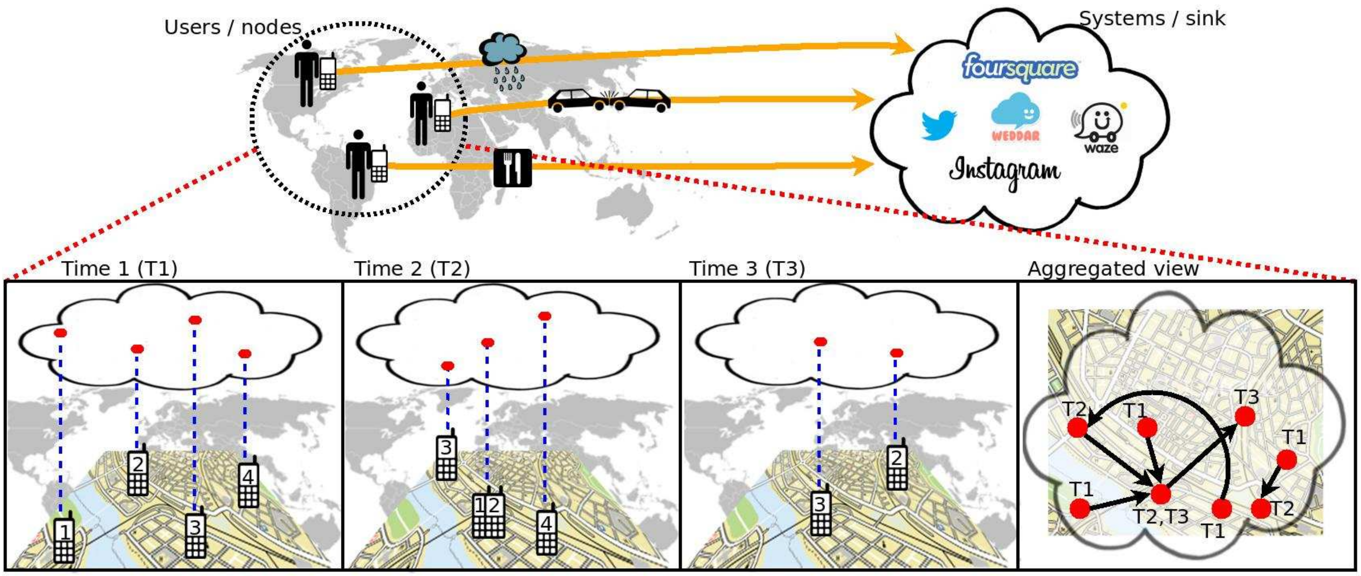

Fig. 1. Illustration of a participatory sensor network.

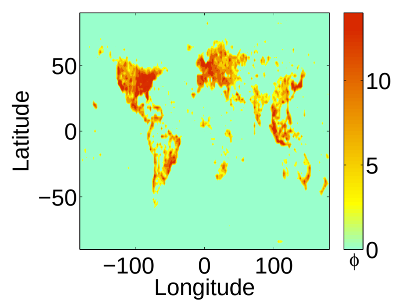

(a) World-wide coverage

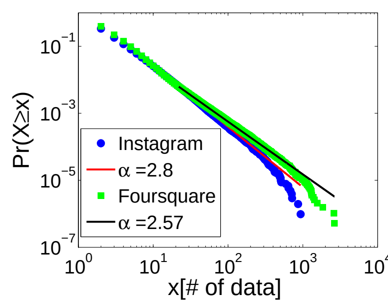

(b) Amount of shared data per locations

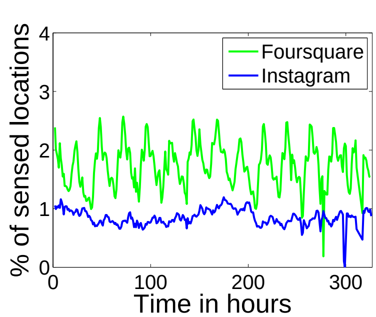

(c) Percentage of sensed locations over time

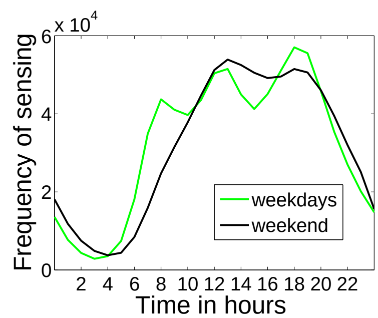

(d) Temporal sensing patterns

Fig. 2. Common characteristics of PSSs.

and comprehensive knowledge about the current patterns and dynamics of a city? Societies are inherently very dynamic, i.e., they change constantly over time, and, as the world gets more and more connected, we believe that these changes tend to be even more frequent. Take, for instance, guide books about large cities involved with the Arab Spring, such as Tunis and Cairo. If they do not capture the changes that came from this period, they are already outdated! Similarly, official census data may quickly become obsolete as such efforts are usually undergone at low frequency (e.g., once every 10 years) due to their high costs.

In contrast, PSNs offer up-to-date views about the locations, opinions, likes and dislikes of their users, and thus have the potential to address the aforementioned questions in near real time and, given their coverage (Figure 2), reaching almost every part of the globe. Next, we elaborate on this potential by presenting various techniques and methods that exploit PSN data to support studies on city dynamics. We conclude the section discussing the main challenges to build and deploy such techniques.

A. Visualizing the invisible image of cities

It is common knowledge that cities are not identical and evolve over time, while habits and routines of their inhabitants are typically distinct. Having said that, can we use a PSN to capture the differences among cities based on common routines of their inhabitants? In [5], we proposed a technique called City Image, which provides a visual summary of the city dynamics based on people movements. This technique exploits urban transition graphs to map the movements of users between locations. An urban transition graph is a directed weighted graph $G(V, E)$, where a node $v_i \in V$ is the category of a specific location (e.g., nightlife) and a direct edge $(i, j) \in E$ marks a transition between two categories. That is, an edge exists from node v_i to node v_j if at least one user shared data in a location categorized by v_j just after sharing data in a location categorized by v_i. Weight $w(i, j)$ of an edge is the total number of transitions that occurred from v_i to v_j. Only consecutive data sharings performed by the same user that occurred within 24 hours, starting at 5:00am, are considered for the sake of transition computation.

City Image is a promising technique to enable a better understanding of the city dynamics, helping the visualization of common routines of their citizens. Each cell in the City Image represents the willingness of a transition from a given category at a given place (vertical axis) to another category (horizontal axis). Red colors represent rejection, blue colors represent favoring, and white represents indifference. We exemplify the City Image technique for two cities - Sao Paulo/Brazil (Figures 3a and 3b) and Kuwait City/Kuwait (Figures 3c and 3d), using the PSN derived from the Foursquare-1 dataset. In both cases, we consider weekdays during the day, which represents the typical period for routines, and weekend during the night, a representative period for (out-of-routine) leisure activities.

First, observe that *office* transitions are more likely to happen on weekdays during the day in both cities, as expected. However, note that the City Images of Sao Paulo and Kuwait city have also significant differences that reflect cultural diversities between both cities. For example, the image representing transitions on weekend nights (Figure 3d) show the lack of favorable transitions considering the category *nightlife* for Kuwait. This is not the case in Sao Paulo (Figure 3b), where the transition *food* $\rightarrow$ *nightlife* is highly favorable to happen, meaning that, in Sao Paulo, people like to eat before clubbing at night. In Kuwait city, instead, people are more favorable to perform the transitions *shop* $\rightarrow$ *food* and *food* $\rightarrow$ *home* on weekend nights.

Techniques to provide easy-to-interpret visualizations of common routines of a city's inhabitants, such as the City Image, are valuable tools to help city planners to better understand the city dynamics and make more effective decisions. They can also be used by IT designers to develop customized software services for different categories of individuals, such as recommendation of popular places to visit for tourists, and recommendation of hot spots for taxi drivers (and other classes of workers). They can also support various sorts of social studies, such as, for example, the investigation of the degree of similarity (in terms of human mobility patterns) between different cities. In particular, the City Image technique enables the clustering of cities by similarity: a City Image is essentially a matrix, and the difference between two matrices could be taken as a measure of the similarity distance between the two cities.

The visualization of city dynamics has also been investigated by other research groups. For example, by using Foursquare data, Cranshaw et al. [6] proposed a model to identify distinct regions of a city that reflect current collective activity patterns, presenting new boundaries for neighborhoods. Long et al. [9] classified venues in a city based on user trajectories, captured in a Foursquare dataset. Their work relies on the assumption that venues that appear together in the trajectories of many users probably are taken as geographic topics, e.g., representing restaurants that people usually go after shopping at a mall. The City Image complements these other efforts by proposing a technique that easies the visualization and comparison of the dynamics of cities based on people's habits and routine.

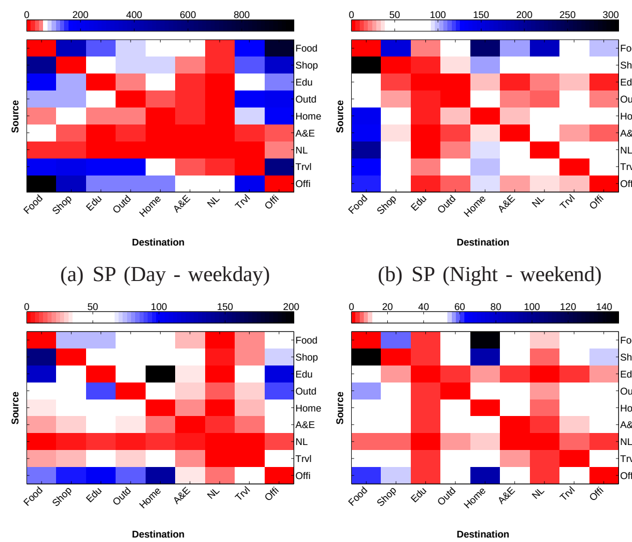

(a) SP (Day - weekday) (b) SP (Night - weekend)

(c) KU (Day - weekday) (d) KU (Night - weekend)

Fig. 3. The *City Image* of Sao Paulo (SP) and Kuwait City (KU) for different periods. Abbreviations of categories of places used in the image (extracted from Foursquare): Arts & Entertainment (A&E); College & Education (Edu); Great Outdoors (Outd); Nightlife Spot (NL); Shop & Service (Shop); and Travel Spot (Trvl).

B. Insights into people movement patterns

Another possible visualization of city dynamics based on data collected by PSNs is illustrated in Figure 4a. It shows a heatmap of the sensing activity for the city of Belo Horizonte (Brazil) for the PSN derived from the Foursquare-1 dataset. The darker the color, the higher the number of check-ins in the area. This heatmap by itself conveys information related to the popularity of specific areas, being thus only partially informative about the city dynamics. Richer information can be obtained by making a small change in the City Image transition graphs presented above (Section III-A). Specifically, we build a graph $G(V, E)$ where node $v_i \in V$ is a given location (e.g., Times Square), instead of a location category. An edge in this new graph represents a transition between two locations, and its weight is the number of people that made that transition. Such modification allows us to draw insights into how people move in the city.

Figure 4b shows the top 50 heavy weighted edges and the top 50 hub nodes (largest degrees) for Belo Horizonte. Stars represent the hubs, black arrows represent edges, and black circles represent self-loops. The larger the symbol, the larger the associated value (edge weight or node degree). Note that the flow of people is very concentrated and skewed, as expected, with a small fraction of the city areas having most of the heavy weighted edges and hubs. Note also that most of the heavy weighted edges are self-loops and short distance edges, implying that people tend to perform activities in their neighborhoods. In contrast, cities that are known for their fast public transportation systems favor the existence of some long distance heavy weighted edges along the public transport links, as we showed in [2].

Note also that, in our modified City Image transition graph, a hub represents a location that people may arrive and depart with high probability. This is a measure of the popularity of different locations in the city that may be exploited by many applications. For example, a public information dissemination scheme, such as the public displays, may benefit from the knowledge of the location of city hubs to make more effective placement and dissemination decisions.

Other related efforts include the work by Cheng el al. [10]: by analyzing check-ins, the authors found that users follow simple and reproducible patterns, and that social status are coupled with mobility. Similarly, Cho et al. [11] investigated human movement patterns and how social ties may impact such movements. They observed that short-ranged travel is both spatially and temporally periodic and is not affected by the social network structure, while long-distance travel is more influenced by social network ties. All these efforts illustrate the increasing interest and the potential of exploiting data shared in PSNs to study human mobility patterns in large scale.

C. Points of interest

Every city has a set of *points of interest* (POIs), that is, areas that attract more attention of inhabitants and visitors than the others. The main touristic sights of a city are included in its set of POIs. However, POIs include more than only the city sights. For example, an area of local pubs in the suburb can be quite popular among city residents, but completely unknown and unpopular among tourists. Thus, an application that naturally emerges from PSNs is related to the identification of POIs in a city. The assumption behind this idea is that each data sharing represents, implicitly, an interest of an individual at a given moment. Thus, the data sharing by many users in a particular location at a given moment serves as evidence that this place is a POI. One advantage of using PSNs to identify POIs in a city is that such technique is robust to dynamic changes. That is, because PSNs sense dynamic data, they can automatically capture the changes in people interests over time, helping to quickly identify areas that suddenly become a POI (e.g., due to the opening of a new restaurant) or that loose popularity.

In [1], we presented a technique to identify POIs and, from them, extract the main sights of a city. The technique considers that each pair i of coordinates (longitude, latitude) $(x, y)_i$ is associated with a point p_i that represents a shared data, e.g. a photo. We start by computing the geographic distance between each pair of points (p_i, p_j) and grouping together the points p_i that are close to each other, e.g., those that have a distance smaller than a certain threshold. Depending on the threshold used, this process produces a very large number of clusters, which are not very informative. To capture the POIs, we use a random model to exclude groups that may have been generated by random situations (i.e., random people movements), and thus do not reflect the dynamics of the city. To identify those groups, we analyze the number of data sharings in each group and adopt simple statistical methods, as described in [1].

To separate sights out of POIs, we first generate a transition graph $G(V, E)$ as described previously, i.e., vertices $v_i \in V$ represent POIs, and there is an edge (i, j) from the vertex v_i to the vertex v_j if a user shared data on a POI v_j after having shared a data 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 by all users. To identify sights, we assume that most tourists follow a well-known path within the city, being guided by its main sights. Moreover, at each POI, he/she shares one or more data and goes to the next tourist spot. Thus, we consider that heavy weighted edges represent frequent transitions from one sight to another, taking their endpoints as the main sights of the city. To obtain these vertices, we exclude from G all edges (i, j) with weights $w(i, j)$ smaller than a threshold w_t, determined based on the probability of randomly generating $w(i, j)$, in a random graph $G_R(V, E_R)$, composed by random edges E_R between POIs $v_i \in V$. In this way we simulate random walks. In possession of the distribution of edge weights, we are able to determine a threshold w_t such that the probability of an edge weight $\geq w_t$ in $G_R(V, E_R)$ is close to zero. All transitions $v_i \rightarrow v_j$ in G with $w(i, j)$ greater than or equal to w_t are assumed to be between sights, according to our conjecture (please refer to [1] for more details).

We applied this technique for the Belo Horizonte city, being able to find the majority of its POIs and sights. The sights identified using the PSN derived from the Instagram-1 dataset, presented in [1], are shown in Figure 4c, whereas Figures 4d and 4e show the sights identified by applying our technique to the PSNs derived from the Instagram-2 and Foursquare-2 datasets, respectively. We note that all PSNs provide useful data to the identification of relevant sights of Belo Horizonte. Indeed, all the eight landmarks recommended by TripAdvisor[6] as the most important cultural and leisure areas of the city are among the identified sights. The figures also show that different PSNs may provide complementary data, as no single PSN found all sights. Such differences may reflect changes in the city during the time when a specific dataset was collected. For example during the collection of Instagram-1, Belo Horizonte was not receiving soccer games. This explains why the soccer stadium was not identified by the PSN derived from that dataset. In contrast, the analysis of a more recently

[6] www.tripadvisor.com.

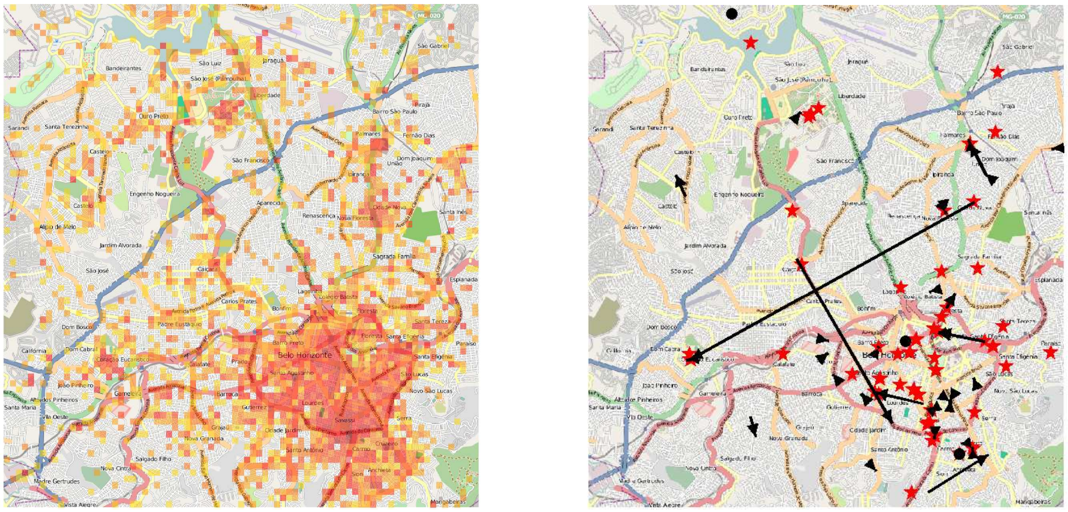

(a) Heatmap

(b) Featured places (nodes) and transitions (edges)

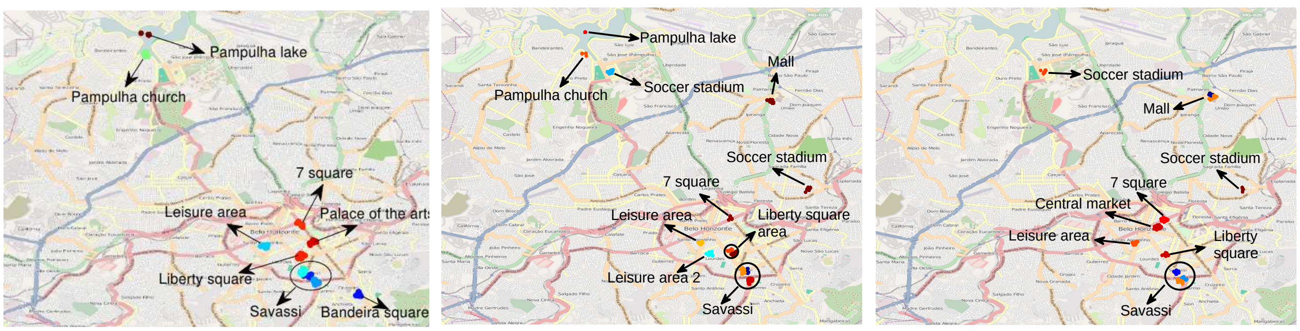

(c) Sights – Instagram-1 (DCOSS [1])

(d) Sights - Instagram-2

(e) Sights – Foursquare-2

Fig. 4. Map views of Belo Horizonte constructed from different participatory sensing data. (a) Heatmap of the number of check-ins, where the color range from yellow (low intensity) to red (high intensity). (b) Top 50 edge weights (transitions between places) and hubs (places that receive people from different venues). Stars represent hubs, black arrows, edges, and black circles, self-loops. (c), (d), and (e) Sights identified in Belo Horizonte using three different datasets.

collected dataset from the same application – Instagram-2– correctly identified the stadium as a major sight of the city. This illustrates how our technique can automatically adapt to changes in the city dynamics, detecting unusual and trendy locations as well as uncovering possibly unexpected patterns. This is thanks to the real time nature of PSNs.

We also note that inherent differences in the applications from which the PSNs are derived may also explain the diversity and complementarity in the set of sights identified by those PSNs. That is because the use of each system might favor the identification of certain types of sights. For example, Figure 4 indicates that the Pampulha Lake and Pampulha Church were identified only by the PSNs derived from Instagram. For these particular places, it is expected that users share more photos than check-ins.

The identification of points of interest in a city was also investigated by Crandall et al. [8]. In their work the authors showed how to organize a large number of geo-tagged photos, jointly analyzing both the textual content of tags and the geospatial data of the photo, with the goal of inferring the location of a photo without using the geospatial data. In the same direction, Kisilevich et al. [12] used geo-tagged photos to analyze and compare temporal events that happened in a city, and also to rank sightseeing places. Our work differs from previous ones since it focus on Instagram, a popular PSS.

D. Socio-economic aspects

Data collected from PSNs can be used to infer the social network topology and dynamics of entire cities, ultimately enabling the analysis of social, economic, and cultural aspects of its inhabitants. For example, one might argue that a small coverage of a certain area by a PSN (i.e., only a small amount of data shared in that area) might indicate a lack of technology access by the local population, since the frequent use of location sharing services often relies on smartphones and 3G or 4G data plans, which, usually, are expensive.

In [2] we showed that the analysis of PSNs enables the visualization of interesting facts related to socio-economic issues of a city. For instance, the PSN derived for the Rio de Janeiro city (Brazil) indicates very small sensing activity in poor areas, including those very close to wealthy ones, a well-known socio-geographic characteristic of that particular city. This information may be useful to drive better public politics in those areas. The same information could be obtained using traditional methods such as surveys, but at the expense of a much slower and expensive process. Towards better understanding social patterns from the analysis of social media, Quercia et al. [13] studied how social media communities resemble real-life ones. The authors tested whether established sociological theories of real-life social networks still hold in Twitter. Lathia et al. [14] examined the correlation between the flow of public transport and census-based indexes of the well-being for the city of London, showing that large scale monitoring of well-being of the communities is possible via the passive sensors used in urban areas.

E. Cultural differences

When studying the social behavior of particular areas, one of the first questions that emerges is: how different is one's culture from another? To address this question, it is necessary to define culture first. However, culture is such a complex concept that no simple definition or measurement can capture it. Among the various aspects that define the culture of a society (or person), one may cite its arts, religious beliefs, letters, and manners. Moreover, eating habits are also fundamental elements in a culture and may significantly mark social differences, boundaries, bonds, and contradictions. Thus, one may use this aspect to study the idiosyncrasies of different societies.

To that end, we consider a PSN derived from Foursquare, and exploit the check-ins given in different types of restaurants, e.g. Japanese, as signals given by users expressing their food preferences. This enables the identification of individual preferences, such as the taste for fish and chips, as well as temporal habits, such as the time and day of the week one goes to a restaurant. This analysis surprisingly says a lot about differences and similarities among cultures. Figure 5a shows the Pearson correlation coefficient between the check-ins given in different types of restaurants for several cities around the world. As we can see, cities in the same country, where inhabitants typically share similar culture and eating habits, have the strongest correlations of restaurant preferences. Besides preferences for food categories, we can also see differences in the periods of time when people go to restaurants and share data. Figures 5b and 5c show the numbers of checkins on restaurants[7] shared by users throughout the hours of the day, on weekdays, in different cities of Brazil and US. These figures capture important differences between the cultures of both countries: while dinner is the main meal for Americans, lunch plays a more important role in the eating habits of Brazilians. This shows that temporal aspects can also give valuable information about cultural differences among cities.

Cultural differences was also studied by Hochman et al. [7], who investigated color preferences in pictures shared through Instagram. The authors found considerable differences across pictures of countries with distinct cultures. In a similar direction, Poblete et al. [15] investigated how tweeting behavior varies across countries, as well as the possible explanations for these differences. The investigation of cultural differences across different cities and countries are valuable in many fields and can support various applications. For example, since culture is an important aspect for economic reasons, the identification of similarities between places that are geographically separated might be valuable for companies that have businesses in one country and want to assess the compatibility of preferences across different markets.

[7]Values are normalized by the maximum value found in any hour for the specific city.

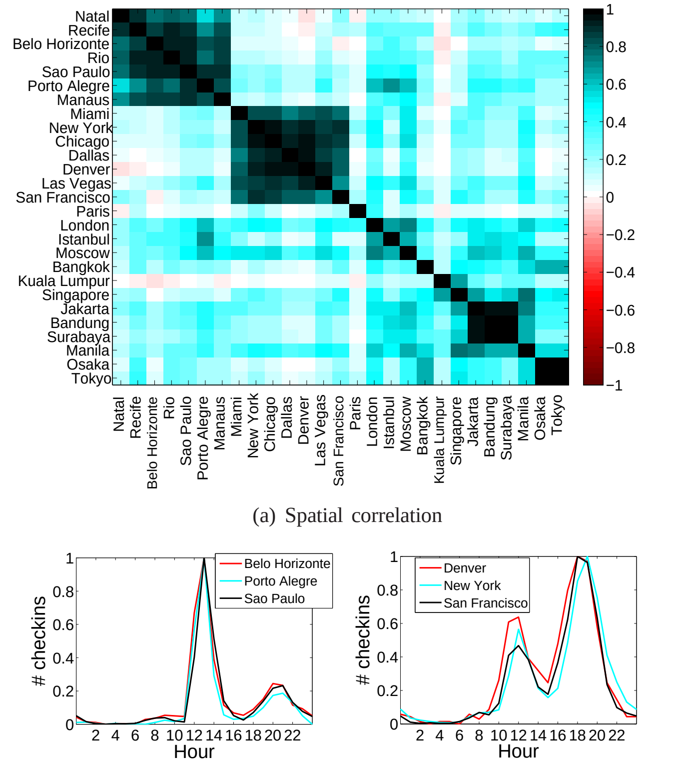

(a) Spatial correlation

(b) Brazilian cities

(c) American cities

Fig. 5. (a) Correlations between restaurant preferences among 27 populous cities. (b) and (c) Average number of check-ins in restaurants during weekdays throughout the hours of the day for 6 different cities.

F. Challenges

In the previous sections we demonstrated the potential of PSNs through different applications and techniques. We now turn to the various challenges one must face to develop such applications and techniques. Some key challenges, which are centered around the data obtained from PSNs, are: **(1) Data quality:** the quality of the shared data when dealing with ubiquitous user contribution is not always guaranteed. Although this issue has been relatively well tackled in the web domain, the ubiquity of contributions might lead to unique problems, for example, sensor readings could be produced by users with little effort [16]. Among other initiatives, Saroiu et al. [16] tackled this issue proposing a trusted platform module, which confirms the integrity of sensing devices; **(2) Data collection:** when dealing with third-parties services, such as Waze, public data might not be available, since users could decide to not publicly disclose it. Incentive mechanisms, such as micro-payments [17], are possible solutions being investigated for this issue; **(3) Big data issues:** PSSs can provide huge volume of data, imposing challenges for *real time* storing, processing, and indexing. The research on big data challenges is very active, and has recently made great advances by, for example, relying on parallel platforms (e.g., Hadoop[8]) for processing large scale datasets.

IV. Discussion and future work

Applications are becoming increasingly mobile, designed to infer user interests and location, and make different sorts of predictions. A mobile device is not just a better option, but may be the only option for many people. This is similar to another phenomenon that has happened for some years now: more and more people are ditching their landlines in favor of cellphones. When people get to that point, they tend to acquire not any cellphone but preferably the latest generation smartphone.

In the previous section, we discussed several applications of PSNs that enable a deeper understanding of a city's dynamics and social behavior. More generally, data from different PSNs can be represented as sensing layers, each one enabling the access of data related to a certain aspect of the city. Figure 6 illustrates this idea, displaying four different layers: **traffic alerts** (obtained, for example, from Waze[9]); **check-ins** (provided, for instance, by Foursquare); **weather condition** (obtained, for example, from Weddar[10]); and **pictures of places** (obtained, for instance, from Instagram). All layers sense data with common features: an identifier of the user who shared the data; the time when the data was shared; and the location where it was shared. The layers differ among themselves by the type of data they provide, difference that we call here *specialty data*. Despite differences in their main purposes, different layers can be used to infer complementary information about the same given city aspect. For example, the traffic alerts layer offers structured data about the traffic, such as information about car accidents or traffic jams. Despite not being the central goal of the pictures of places layer, conditions about the traffic could also be inferred by analyzing the content of the shared pictures that were taken on a given road.

We here discussed current and previous efforts to explore PSNs derived from two types of systems, namely location and photo sharing applications. These systems fall into two different sensing layers: check-ins and picture of places. We have also illustrated the potential of obtaining enriched and more complete information by jointly analyzing both layers (e.g., in the POI detection technique). We particularly have studied the time and location dimensions of our data from both layers. Besides that, we have explored also the specialty data provided by the check-ins layer (e.g., in the City Image technique).

Certainly, a range of fruitful opportunities may emerge when exploring the specialty data offered by other layers. For example, near real-time traffic information provided by the traffic alerts layer could be leveraged to improve algorithms for navigation services. Car accidents, potholes, and slippery roads are examples of valuable pieces of information whose detection is hard to obtain with traditional sensors, but becomes more feasible with the use of the traffic alerts layer. This is possible because each sensing layer is composed by crowdsourced data, providing information on a personal level. It is worth noting also that since each layer represents a partial view of the city, their aggregation can provide a deeper understanding of it.

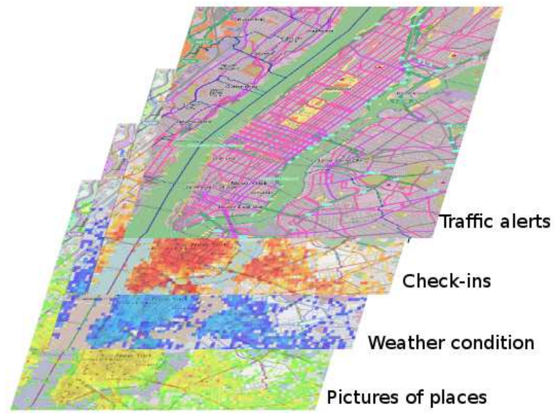

Fig. 6. Sensing layers for the city of New York. Each layer gives information about a specific aspect of the city. Data from different layers may be aggregated in order to better understand a given context or to answer a given question, e.g., "give me a sunny and crowded location where I can take a good picture and that I can easily arrive by car."

In summary, the use of PSNs can help us better understand the dynamics of cities, and from this understanding we are able to offer smarter services to meet people's needs. A future direction we intend to pursue is jointly analyzing different PSNs (sensing layers). We also intend to investigate the interplay between data obtained from traditional wireless sensor networks and data obtained from PSNs.

[8] http://hadoop.apache.org.
[9] http://www.waze.com.
[10] http://www.weddar.com.

References

[1] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, J. Salles, and A. A. F. Loureiro, "A picture of Instagram is worth more than a thousand words: Workload characterization and application," in *Proc. of IEEE DCOSS'13*, Cambridge, USA, 2013, pp. 123–132.

[2] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, and A. A. F. Loureiro, "Challenges and opportunities on the large scale study of city dynamics using participatory sensing," in *Proc. of IEEE ISCC'13*, Split, Croatia, 2013.

[3] J. P. Scott and P. J. Carrington, *The SAGE Handbook of Social Network Analysis*. Sage Publications Ltd., 2011.

[4] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil, "An empirical study of geographic user activity patterns in Foursquare," in *Proc. of ICWSM'11*, Barcelona, Spain, 2011, pp. 570–573.

[5] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, J. Salles, and A. A. F. Loureiro, "Visualizing the invisible image of cities," in *Proc. of IEEE CPSCom'12*, Besancon, France, 2012.

[6] 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. of ICWSM'12*, Dublin, Ireland, 2012, pp. 58–65.

[7] N. Hochman and R. Schwartz, "Visualizing instagram: Tracing cultural visual rhythms," in *Proc. of the Workshop on Social Media Visualization*, Dublin, Ireland, 2012.

[8] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg, "Mapping the world's photos," in *Proc. of ACM WWW'09*, Madrid, Spain, 2009, pp. 761–770.

[9] X. Long, L. Jin, and J. Joshi, "Exploring trajectory-driven local geographic topics in foursquare," in *Proc. of ACM Ubicomp'12*, Pittsburg, USA, 2012, pp. 927–934.

[10] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui, "Exploring millions of footprints in location sharing services," in *Proc. of ICWSM'11*, Barcelona, Spain, 2011, pp. 81–88.

[11] E. Cho, S. A. Myers, and J. Leskovec, “Friendship and mobility: User movement in location-based social networks,” in *Proc. of KDD’11*, San Diego, USA, 2011, pp. 1082–1090.

[12] S. Kisilevich, M. Krstajic, D. Keim, N. Andrienko, and G. Andrienko, “Event-based analysis of people’s activities and behavior using flickr and panoramio geotagged photo collections,” in *Proc. of the 14th Int’l Conf. Information Visualisation*, Washington, DC, USA, 2010, pp. 289–296.

[13] D. Quercia, L. Capra, and J. Crowcroft, “The social world of Twitter: Topics, geography, and emotions,” in *Proc. of ICWSM’12*, Dublin, Ireland, 2012, pp. 298–305.

[14] N. Lathia, D. Quercia, and J. Crowcroft, “The hidden image of the city: sensing community well-being from urban mobility,” in *Proc. of Pervasive’12*, Newcastle, UK, 2012, pp. 91–98.

[15] B. Poblete, R. Garcia, M. Mendoza, and A. Jaimes, “Do all birds tweet the same?: Characterizing Twitter around the world,” in *Proc. of the 20th ACM CIKM’11*, Glasgow, UK, 2011, pp. 1025–1030.

[16] S. Saroiu and A. Wolman, “I am a sensor, and i approve this message,” in *Proc. of ACM HotMobile’10*, Annapolis, Maryland, 2010, pp. 37–42.

[17] S. Reddy, D. Estrin, M. Hansen, and M. Srivastava, “Examining micro-payments for participatory sensing data collections,” in *Proc. of ACM Ubicomp’10*, Copenhagen, Denmark, 2010, pp. 33–36.

Thiago H. Silva graduated with a BSc in Computer Science in 2004. In 2009 he obtained his MSc in Computer Science from the Federal University of Minas Gerais (UFMG), where he is a PhD student. Thiago was a research intern at Telecom Italia, Venezia-Italy, and a visiting PhD student at the University of Birmingham, Birmingham-UK, and at INRIA, Paris-France. Some reseach areas of interest: Ubiquitous computing; urban computing; social computing; and user behavior modeling.

Pedro O.S. Vaz De Melo is an assistant professor in the Computer Science Department of Federal University of Minas Gerais (UFMG). He has degree (2003) and Masters (2007) in Computer Science from the Pontifical Catholic University of Minas Gerais (2003). He got his Ph.D. at Federal University of Minas Gerais (UFMG) with a one year period as a visiting researcher in Carnegie Mellon University and also spent five months as a visiting researcher at INRIA Lyon.

Jussara M. Almeida has a Bachelor and Master degrees from the Computer Science Department of the Universidade Federal de Minas Gerais (UFMG), in Brazil. She also has an M.Sc. and a Ph.D. degrees from the University of Wisconsin-Madison, US. She is currently an Associate Professor of Computer Science at UFMG. Her research interests are primarily on performance modeling of large-scale distributed systems, workload and user behavior modeling, and social computing.

Antonio A.F. Loureiro received his B.Sc. and M.Sc. degrees in computer science from the Federal University of Minas Gerais (UFMG), Brazil, and the Ph.D. degree in computer science from the University of British Columbia, Canada. Currently, he is a full professor of computer science at UFMG, where he leads the research group in ubiquitous computing. His main research areas are wireless sensor networks, ubiquitous computing, and distributed algorithms.