Urban Sensing Using Mobile Phone Network Data: A Survey of Research

FRANCESCO CALABRESE, IBM Research-Ireland, Dublin, Ireland
LAURA FERRARI, Dipartimento di Scienze e Metodi dell'Ingegneria, University of Modena and Reggio Emilia, Italy
VINCENT D. BLONDEL, Université Catholique de Louvain, Louvain-la-Neuve, Belgium

The recent development of telecommunication networks is producing an unprecedented wealth of information and, as a consequence, an increasing interest in analyzing such data both from telecoms and from other stakeholders' points of view. In particular, mobile phone datasets offer access to insights into urban dynamics and human activities at an unprecedented scale and level of detail, representing a huge opportunity for research and real-world applications. This article surveys the new ideas and techniques related to the use of telecommunication data for urban sensing. We outline the data that can be collected from telecommunication networks as well as their strengths and weaknesses with a particular focus on urban sensing. We survey existing filtering and processing techniques to extract insights from this data and summarize them to provide recommendations on which datasets and techniques to use for specific urban sensing applications. Finally, we discuss a number of challenges and open research areas currently being faced in this field. We strongly believe the material and recommendations presented here will become increasingly important as mobile phone network datasets are becoming more accessible to the research community.

Categories and Subject Descriptors: C.2.0 [**Computer-Communication Networks**]: General—*Data communications*; I.1.4 [**Symbolic and Algebraic Manipulation**]: Applications; I.6.3 [**Simulation and Modeling**]: Applications; J.1 [**Computer Applications**]: Administrative Data Processing; K.4.1 [**Computers and Society**]: Public Policy Issues

General Terms: Algorithms, Experimentation, Measurement

Additional Key Words and Phrases: Mobile phone data, data mining, urban planning, urban transportation

ACM Reference Format:
Francesco Calabrese, Laura Ferrari, and Vincent D. Blondel. 2014. Urban sensing using mobile phone network data: A survey of research. ACM Comput. Surv. 47, 2, Article 25 (November 2014), 20 pages.
DOI: http://dx.doi.org/10.1145/2655691

1. INTRODUCTION

Over the past decade, the development of digital networks has produced an unprecedented wealth of information reflecting various aspects of urban life. These digital traces are valuable sources of data in capturing the pulse of the city in an astonishing degree of temporal and spatial detail and could be used to make urban systems more efficient.

Authors' addresses: F. Calabrese, IBM Research-Ireland, Damastown Industrial Park, Mulhuddart, Dublin 15, Ireland; email: fcalabre@ie.ibm.com; L. Ferrari (current address), via Borgo Costa 16, 42027 Montecchio Emilia, Reggio Emilia, Italy; email: lauraferrari85@gmail.com; V. D. Blondel, Avenue Georges Lemaître 4, boîte L4.05.01, B-1348 Louvain-la-Neuve, Belgium; email: vincent.blondel@uclouvain.be.

DOI: http://dx.doi.org/10.1145/2655691

The International Telecommunication Union estimates that at the end of 2011, there were 6 billion mobile subscriptions, with a global penetration of 87%, and 79% in the developing world [ITU 2011]. Every mobile phone leaves digital traces while interacting with its infrastructure. Thus, each phone can be seen as a mobile sensor that allows one to detect the geographic position of the subscriber holder almost in real time. Telecom operators are aware of the potential of such data and they have recently started to experiment with new business models in which they would generate revenues not only from their final customers (mobile phone users) but also from upstream customers such as traffic analysis, social networking, and advertising companies. As a result, they have started sharing aggregated mobile data with various research communities [Technology Review 2010]. Thanks to that, massive datasets about cell phone users have been exploited in a variety of urban-related applications, including understanding mobility patterns [González et al. 2008; Isaacman et al. 2010], the use of urban spaces [Reades et al. 2007], travel demand during special events [Calabrese et al. 2010], social network structure [Onnela et al. 2007], and geographical dispersal of mobile communications [Lambiotte et al. 2008]. More recently, research challenges have also been proposed by Orange,[1] Telefonica,[2] and Telecom Italia,[3] where operators have released telecommunication data to the wide research community, which are now accessible and studied by hundreds of research laboratories around the world. Clearly, using mobile phone data for urban sensing could have a great impact in developing countries, where specific sensors (such as traffic sensors) are rarely put in place. A recent primer from the UN Global Pulse organization summarizes the latest research examples addressing developing countries.[4]

While several research works have been conducted on using different types of mobile phone network data for specific purposes, each work has been done on a specific flavor of the data (different accuracy, granularity, aggregation level), and so it is difficult to understand whether a particular technique could indeed be applied to a different dataset and what results that would provide. At the same time, if a researcher or practitioner is interested in building a specific urban sensing application, it is difficult for him or her to figure out which particular mobile phone network dataset would be the most suitable and which techniques should be applied to the data to achieve the specific goal. This article surveys the new ideas and techniques related to the use of telecommunication data for urban sensing, with the specific goal to help researchers and practitioners navigate the variety of mobile phone network datasets and associated processing techniques that have been presented in the literature to build urban sensing applications. More specifically, Section 2 shows what telecom data can tell about urban dynamics. Section 3 outlines the mechanisms at the basis of mobile phone data generation. Section 4 surveys the filtering and processing techniques proposed so far to extract insights from this data and summarizes them to provide recommendations on which datasets and techniques to use for specific applications. Finally, Section 5 provides an overview of the challenges currently being faced in this field, and Section 6 concludes.

2. MOBILE PHONE NETWORK DATA FOR URBAN ANALYSIS

It is well known that 50% of the globe's population lives in urban areas, which cover only 0.4% of the earth's surface [Fund 2007], and 70% are projected to do so by 2050. From one side, such urbanization opens great opportunities for improving people's

[1]http://www.d4d.orange.com.
[2]http://dynamicinsights.telefonica.com/674/the-details.
[3]http://www.telecomitalia.com/tit/en/bigdatachallenge.html.
[4]http://www.unglobalpulse.org/Mobile_Phone_Network_Data-for-Dev.

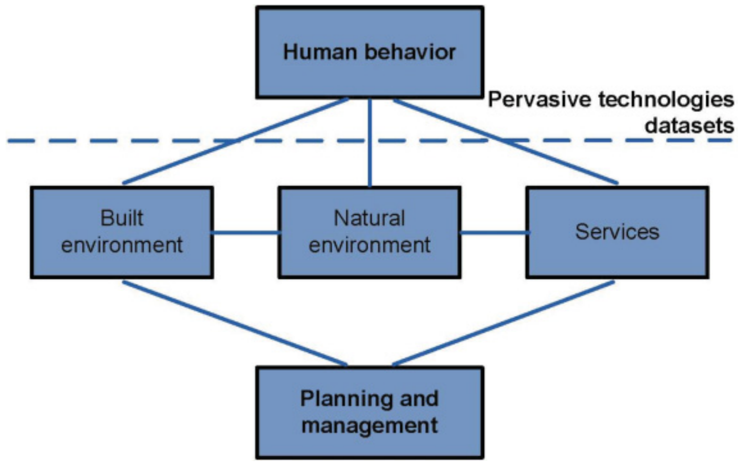

Fig. 1. Schema reflecting the role of pervasive technologies datasets in an urban scenario.

lifestyles; from the other side, there is the need to prevent a potential economic, health, and environmental disaster [Manyika et al. 2011]. Pervasive technologies datasets are a way to understand how people use the city's infrastructure from the point of view of mobility (e.g., transportation mode), consumption (e.g., energy, water, waste), and environmental impact (e.g., noise, pollution). In fact, this kind of information offers new insights about the city (see, for example, the *Villevivante* project[5]), which are of great interest both from an economic and from a political perspective. In particular, urban planning can benefit from the analysis of personal location data. Decisions that can be improved by analyzing such data include the mitigation of traffic congestion and planning for high-density development. Urban transit and development planners will increasingly have access to a large amount of information about peak and off-peak traffic hotspots, volumes, and patterns of transit use with which they can potentially cut congestion and the emission of pollutants. By drilling down into this wealth of data, urban planners will be more informed when they make decisions on anything from the placing and sequencing of traffic lights to the likely need for parking spaces. As an example, Singapore's public transportation is already using 10-year demand forecasts partly based on personal location data to plan transit needs[6] and is continuing to invest in this direction through the Future Urban Mobility initiative.[7] Figure 1 shows how pervasive technologies datasets fit in this scenario. The human behavior of people in a city reflects how citizens use the built environment, the natural environment, and the services offered by a city. Pervasive technologies are able to capture human behaviors and produce related datasets that contain very useful information for planning and management.

One important pioneering work in the field of community dynamics sensing using cell phone data has been conducted within the *Reality Mining* project.[8] Reality mining deals with the collection and analysis of machine-sensed environmental data pertaining to human social behavior, with the goal of identifying predictable patterns of behavior. Mobile phones (and similarly innocuous devices) are used for data collection, opening social network analysis to new methods of empirical stochastic modeling [More and

[5]http://villevivante.ch.
[6]http://www.onemotoring.com.sg/publish/onemotoring/en/on_the_roads/traffic_management.html.
[7]http://smart.mit.edu/research/future-urban-mobility/future-urban-mobility.html.
[8]http://realitycommons.media.mit.edu/.

Lingam 2013]. The Reality Mining project collected data by asking volunteers to carry cell phones programmed to measure and store sensor data. In this survey, instead, we focus on data opportunistically collected by the telecom operators by product of their own operations, without the requirements of people to carry specific devices or agree to install or enable specific features on their phones. Of course, privacy is a real issue in using this kind of technology. In fact, every country has its own regulations that telecommunication operators have to comply with. The main worry arising from the use of mobile phone network data is the fact that phone users' movements are monitored, particularly in cases where such personal location data are made available to applications whose beneficiaries are third parties. As an example, the European Directive 2002/58/EC regulates the treatment of personal data and protection of intimacy in the electronic communications sector.[9] Article 14 of this directive includes a description of location data, stating that "Location data may refer [...] to the identification of the cell in the network in which the mobile terminal is located at a given moment or to the time at which the localization information has been registered." Article 9 of this directive also supplies regulations covering location data, as follows: "In the event that location data can be processed [...] such data may only be processed if they are made anonymous, or with the prior consent of the users or clients, to the extent and for the time necessary to provide a value-added service." Thus, in order to be compliant with regulations, all the data used for the research in this field (see the list of references) has been released by telecom operators so that it was impossible to associate the location data with actual cell phone numbers.

In the field of urban analysis, mobile phone network data has been used in several research efforts:

(1) **Estimating population distribution.** With this regard, the use of mobile phone network data is twofold: (1) estimate where people live and (2) estimate how population density changes over time, that is, identify regions densely populated during particular days of the week and hours of the day. In particular, from one side, the focus is on identifying locations that are meaningful to users. Ahas et al. [2010] and Isaacman et al. [2011] introduce a model for determining the geographical location of home and workplaces, while Nurmi and Bhattacharya [2008] describe and evaluate a nonparametric Bayesian approach for identifying places from sparse GPS traces (given the generic approach of the methodology, it can be easily applied to mobile phone network data). From the other side, the focus is on analyzing how the density of people changes over time. For example, Sohn et al. [2006], Sevtsuk and Ratti [2010], and de Jonge et al. [2012] explore how coarse-grained GSM data from mobile phones can be used to recognize high-level properties of user mobility and daily step count. The work in Krisp [2010] shows how calculating and visualizing mobile phone density can assist fire and rescue services. Moreover, in Soto et al. [2011], the information derived from the aggregated use of cell phone records is used to identify the socioeconomic levels of a population.

(2) **Estimating types of activities in different parts in the city.** During the week, the call activity of a residential region, a commercial region, or a business is different. It may be possible to derive a classification from the call activity profile of a region, thus allowing one to classify regions as "residential," "commercial," or "business." For example, Girardin et al. [2009] provide a case study where aggregate and anonymous cell phone network activity data and georeferenced photos from Flickr

[9]Directive 2002/58/EC of the European Parliament and of the Council of July 12, 2002, concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on Privacy and Electronic Communications). http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:en:HTML.

allow one to track the evolution of the attractiveness of different areas of interest in New York. Reades et al. [2007] monitor the dynamics of Rome and obtain clusters of geographical areas measuring cell phone tower activity. Other works try to focus on the specific land use of a city. For example, Soto and Frias-Martinez [2011] use time series analysis to automatically identify land uses from aggregated call detail record databases. The work is focused on the following types: industrial parks and office areas, commercial and business areas, nightlife areas, leisure and transport hubs, and residential areas.

(3) **Estimating mobility patterns.** Using the cell phone ID, timestamp, and location data of an event (call, SMS, Internet usage), it is possible to estimate commuters' mobility in predefined regions. Several groups of researchers did extensive work in this field. To name a few, the *Barabasi Lab*[10] has an open project on "Individual Mobility Patterns." For example, González et al. [2008] show how the widespread coverage of mobile phone wireless networks in urban areas makes it possible to track both groups and individuals. Song et al. [2010] investigated to what degree human behavior is predictable with results indicating that the development of accurate predictive models is a scientifically grounded possibility, with potential impact on our well-being and public health. Moreover, they analyzed several aspects of mobility patterns ranging from human trajectories [Song et al. 2010] to migration [Simini et al. 2012] and road usage patterns [Wang et al. 2012]. Several other important works in this area have been conducted by the *MIT Senseable city lab.*[11] Their works aim to investigate and anticipate how digital technologies are changing the way people live and their implications at the urban scale. In particular, in their first works, the authors used the real-time data collected from mobile phones to monitor the vehicular traffic status and the movements of pedestrians in Rome, Italy [Calabrese et al. 2011]. Finally, Becker et al. [2013] characterized human mobility in several U.S. cities to offer insight on a variety of important social issues such as evaluating the effect of human travel on the environment.

(4) **Analyzing local events.** The increasing availability of mobile phone usage datasets in recent years has led to a number of studies also related to local events and the interplay with mobility. In particular, several works tried to infer the human patterns of mobility during emergencies and special events [Bagrow et al. 2011; Calabrese et al. 2010; Ferrari et al. 2012; Lu et al. 2012; Traag et al. 2011].

(5) **Analyzing the geography of social networks.** The impact of geography on social interactions has been exploited from a statistical perspective [Lambiotte et al. 2008] to derive a geography of mobile communications based on the relative frequency of communications as well as their average duration [Blondel et al. 2010] and to study the social radius of influence at both the communication and mobility scales [Calabrese et al. 2011].

Mobile phone network data has been used not only in research works but also in running products based on both aggregated and individual data. A first group of applications deals with the issue of using mobile phone network data to derive urban traffic patterns. Traditional companies (such as Inrix[12] and Delcan[13]) use traffic collection methods based on locating GPS-enabled vehicles and mobile devices. The use of mobile phone network data in order to leverage traffic information enables one to handle more data nodes (given the huge number of mobile phone subscribers), and

[10]http://www.barabasilab.com.
[11]http://senseable.mit.edu.
[12]http://www.inrix.com.
[13]http://delcantechnologies.com.

therefore offers higher resolution than traditional traffic collection methods that are based on a relatively small group of GPS-enabled vehicles. Thus, an increasing number of telecom operators are making partnerships with external companies that can provide real-time services using traffic information; see, for instance, the partnership between Vodafone and TomTom.[14] For example, Cellint[15] provides a worldwide service using mobile signaling data to locate the cars on the road. Such data is then analyzed to provide immediate incident detection (such as road sensors), as well as travel time and local speed over short segments (e.g., 200 meters in urban areas and 500 meters in other areas) for all the roads within a covered area. Intellimec is a similar company[16] that provides real-time traffic and incident information in Canada. Another company that leverages mobile phone network data to provide traffic information is Airsage,[17] which aggregates signaling data from cellular networks to provide real-time speed and travel times for major roads. The company currently provides real-time location and traffic data in almost every city in the United States. Airsage also tries to provide insight into the behavior of consumers at specific locations and at different times during the day. A similar approach has been taken by Telefonica with the Smart Steps product,[18] which uses anonymized and aggregated mobile network data to provide insights representative of the total population in each area and time.

Other applications focus on using mobile phone network data to provide services based on a more "social" aspect. For example, Sense Networks[19] is commercializing Macrosense, a machine-learning technology model that aggregates historical and real-time mobile phone location data to, for instance, identify the best street corners from which to hail a taxi. Sense Networks' first application for consumers was CitySense, a tool designed to answer the question "Where is everyone going right now?" CitySense showed the overall activity level of the city, hotspots and places with unexpectedly high activity, all in real time. The tool also uses Yelp and Google to show what venues are operating at those locations. CabSense, another Sense Network application realized in early 2010, offers users an aggregated map generated by analyzing tens of millions of data points that rank street corners by the number of taxicabs picking up passengers every hour or every day of the week.

These examples show how mobile phone network data has the potential to do the following: (1) offer the possibility to study micro- and macro-behaviors and (2) truly reflect human behavior given the fact that data is becoming more and more available thanks to the increasing adoption of mobile technologies. The big issue shared by all these works is to validate the extracted insights. To this regard, comparative datasets are useful to:

(1) Validate findings extracted from analysis of the mobile phone network data;
(2) Define scaling factors to extend results to the overall population;
(3) Augment information about urban space, which is useful to extract higher-level patterns.

Table I outlines the main comparative datasets commonly used to validate the results obtained from mobile phone network data and highlights their pros and cons.

In particular:

[14]http://enterprise.vodafone.com/insight_news/case-study/tomtom.jsp.
[15]http://www.cellint.com.
[16]http://www.intellimec.com.
[17]http://www.airsage.com.
[18]http://dynamicinsights.telefonica.com/488/smart-steps.
[19]http://www.sensenetworks.com.

Table I. Pros and Cons of the Main Comparative Datasets

Type	Pros	Cons
Census and surveys	Very refined spatial resolution	Often outdated
Land use	Different categories	Different spatial units
Points of interest	Very refined categories	Different sources of data may provide different categories for the same points of interest

Census and Surveys. Census and surveys provide datasets related to very different areas: demographics, health, education, government and security, communication and transport, and so forth (see, for example, the 2010 U.S. Census[20]). Such datasets can be used to (1) validate home and working areas; (2) validate city patterns such as hotspots, commuting, and traffic flows; and (3) validate land use. The main advantage of this kind of data is the very refined spatial resolution that is often the census block. The main disadvantages are that they are updated usually only every 5 to 10 years. Moreover, only some questions are asked, thus providing only a partial view of human behavior.

Land Use. Global land use datasets offer access to a number of datasets that characterize an area based on its planned use (e.g., the NASA Global Land Use Datasets[21]). Different categories have been defined such as country codes, population density, cultivation intensity, and so forth. The main disadvantages are the possibly different spatial units in which they are aggregated.

Points of Interest. Points of interest are a list of businesses and important places to visit in a city. Usually every point of interest is characterized by a category and a location. There are many possible different sources—Yellow Pages, Yelp, Google Places, and so forth—that might provide different information. As an example, the “A60,” a famous rooftop bar in Manhattan, can be categorized as “Bar” by one source and as “Nightlife” by another source. In most comparisons, categories are aggregated in super-categories (e.g., bars and restaurants are aggregated in the super-category “Food”).

There are some challenges and limitations in comparing different datasets. The main one is that different collection periods and different spatial units introduce difficulties in comparing datasets. For example, census data is aggregated at the block, track, or country level, while mobile phone network data is aggregated at the cell tower level.

Finally, another limitation in the use of mobile phone data to estimate urban dynamics is the potential biases in differential ownership of phones among different demographic groups. A recent study, however [Wesolowski et al. 2013], has shown that for the purpose of estimating human mobility, mobile phone data from a large telecom operator in Kenya seemed robust to biases in phone ownership across different geographical and socioeconomic groups. While this study does not automatically generalize to any mobile phone network dataset, it shows that for large enough samples, the biases have a low impact on the extracted mobility patterns.

In the next section, we will discuss how telecommunication networks generate the mobile phone datasets and their features.

3. MOBILE PHONE NETWORK DATA GENERATION

When a mobile phone is switched on, it regularly notifies of its position in terms of the actual cell where it is currently located. The notification of the mobile phone position can be triggered by *events* (call, SMS, or Internet usage) or by updates of the *network*

[20]http://2010.census.gov.
[21]http://data.giss.nasa.gov/landuse/.

Table II. Example of a CDR Log: Anonymized Originating and Terminating User ID, Originating and Terminating Cell ID, Timestamp, and Call Duration

Originating_id	Originating_cell_id	Terminating_id	Terminating_cell_id	Timestamp	Duration
24393943	10121	17007171	10121	24031517	29
24393943	5621	17007171	2721	25141136	38
24393943	17221	17007171	2521	25534630	188
24393943	31041	17007171	5111	32440483	111
24393943	10121	17007171	9411	33152308	145
24393943	6321	17007171	20921	33431903	132
24393943	7041	17007171	10021	33435718	17
24393943	7021	17007171	14321	34160370	53

Table III. Cell Location Information

Cell id	Lat	Lon
10121	44.658885	10.925102
17221	44.701606	10.628872

(for a more detailed description of the technologies and standards used to derive the position of mobile phones, see Wang et al. [2008]).

Event-Driven Mobile Phone Network Data. Today, there are two primary sources of these data: communication and Internet usage. Most telephone networks generate call detail records (CDRs), which are data records produced by a telephone exchange documenting the details of a phone call or SMS passed through the device. A CDR is composed of data fields that describe the telecommunication transaction, such as the user ID of the subscriber originating the transaction, the user ID receiving the transaction, the transaction duration (for calls), the transaction type (voice or SMS), and so on. Each telecommunication operator decides which information is emitted and how it is formatted. As an example, there could be the timestamp of the end of the call instead of the duration. Table II shows an example of a CDR log, while Table III shows the mapping between cell IDs and locations.

The second source of data is Internet usage. In telecommunications, an IP detail record (IPDR) provides information about Internet Protocol (IP)-based service usage and other activities. The content of the IPDR is determined by the service provider, the Network/Service Element vendor, or any other community of users with authority for specifying the particulars of IP-based services in a given context. Examples of IPDR data fields are user ID, type of website, time of event, number of bytes transmitted, and so forth. It is important to note that the margin of error in this case varies widely according to whether the device to which the IP address is attached is mobile, and to the density and topology of the underlying IP network.

Both communication and Internet usage can be associated to the cell phone towers used during the interaction.

Network-Driven Mobile Phone Network Data. A cellular network is a radio network of individual cells, known as base stations. Each base station covers a small geographical area that is part of a uniquely identified location area. By integrating the coverage of each of these base stations, a cellular network provides a radio coverage over a much wider area. A group of base stations is named a location area (LA), or a routing area. An LA is a set of base stations that are grouped together to optimize signaling (see Figure 2(a)).

Typically, tens or even hundreds of base stations share a single base station controller (BSC). The BSC handles allocation of radio channels, receives measurements from the mobile phones, and controls handovers from base station to base station.

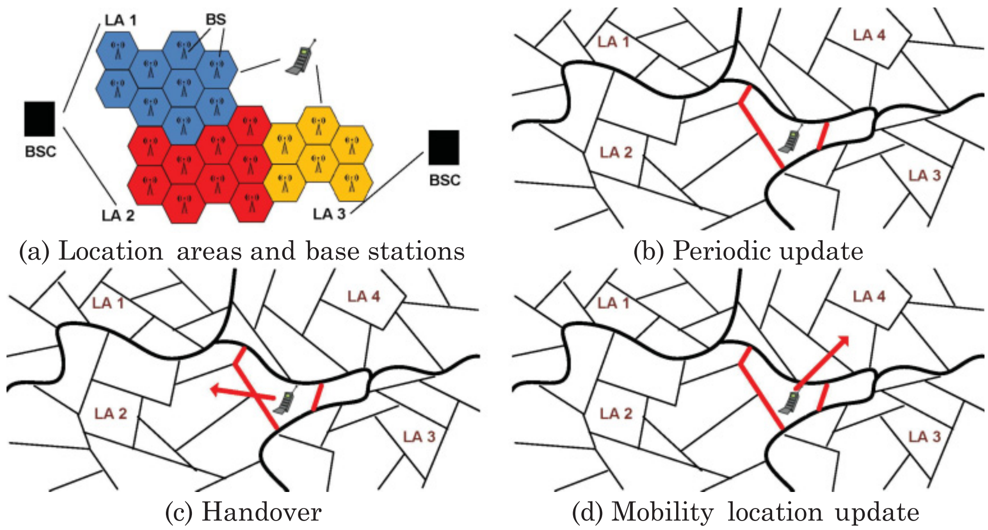

(a) Location areas and base stations (b) Periodic update

(c) Handover (d) Mobility location update

Fig. 2. (a) Location area and base stations. (b) Periodic update. (c) Handover. (d) Mobility location update.

In such a context, different types of location updates can happen:

(1) **Periodic update**, which is generated on a periodic basis and provides information on which cell tower the phone is connected to (see Figure 2(b)).
(2) **Handover**, which is generated when a phone involved in a call moves between two cell areas (see Figure 2(c)).
(3) **Mobility location update**, which is generated when the phone moves between two location areas (see Figure 2(d)).

Location updates also happen when the phone changes the type of connectivity it uses to access the telecommunication infrastructure (e.g., from 2G to 3G). The frequency of these updates strongly depends on how the operator has deployed the different connectivity technologies.

Another important aspect is how the user's location can be detected. Location information can be extracted as part of the interaction data between the mobile phone and the telecommunication infrastructure. In most cases, it is represented by the cell tower position or the cell sector to which the mobile phone is connected. Table II shows an example of a CDR's location information, represented by the *cell id* field. Table III maps each *cell id* to the corresponding latitude and longitude coordinates.

In particular, triangulated location can be estimated as having access to data collected at lower levels in the network. The format of such data is given by standard documentation provided by network operators (see the 3gpp standard documentation[22]). The principal techniques are the following:

(1) **Timing Advance (TA)**, which is a value that corresponds to the length of time a signal takes to reach the cell tower from a mobile phone. Since the users are at various distances from the cell tower and radiowaves travel at the finite speed of light, the precise arrival time can be used by the cell tower to determine the distance to the mobile phone (see Figure 3(a)).
(2) **Received Signal Strength (RSS)**, which is a measurement of the power present in the signal received by cell towers surrounding the phone. Because the power

[22]http://www.3gpp.org/.

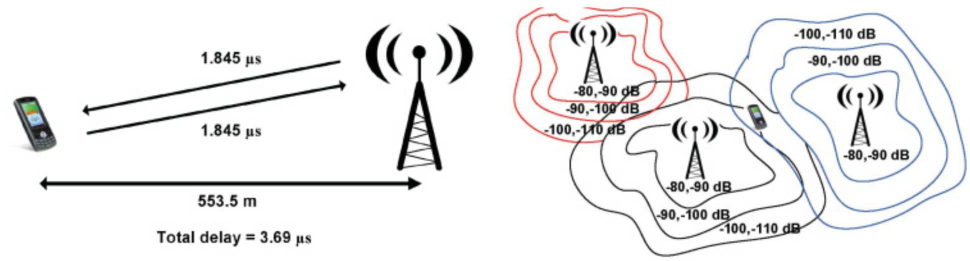

(a) (b)

Fig. 3. Estimating the mobile phone location information: (a) Time Advance and (b) Received Signal Strength techniques.

Table IV. Example of Cell Tower Location Information Obtained Using Propagation Models

User Hash	Longitude	Latitude	Uncertainty	Timestamp
4ba232e4d96f47dc94f7441e87c164fb	16	81	56	1246759931
4ba232e4d96f47dc94f7441e87c164fb	06	09	252	1246759922
4ba232e4d96f47dc94f7441e87c164fb	99	95	208	1246760034

levels at the start of the signal transmission are well known and the power drop in signal in open spaces is well defined, RSS can be used to estimate the distance between a mobile phone and the surrounding cell towers (see Figure 3(b)).

It is important to note that with these methodologies, the accuracy of the mobile phone position is around 500m in urban areas. An accuracy of 150m in urban areas can be obtained using propagation models and irradiation diagrams; such techniques estimate the mobile phone position by finding the point that minimizes the mean square error between measured and estimated mean power received by all base stations. Table IV shows an example of the cell tower location information obtained using propagation models and irradiation diagrams; the main difference is represented by the uncertainty field that gives an estimation of the accuracy (for instance, in meters) of the mobile phone position.

Service providers in each country have different rules and restrictions as to what kind of data can be exchanged through their network. Individual data is rarely available in real time even for service providers. Moreover, the use of individual data can lead to privacy concerns (as explained in Section 2). The same data can be aggregated at different spatial and temporal scales. For example, mobile phone network data can be aggregated at the cell tower level by considering the number of calls, Erlang (total communication time; see Freeman [2005]), the number of SMS, the number of handovers, the number of location updates, and so forth.

Aggregated data can be more easily accessible in real time or with low delay. Moreover, regarding the data volume, aggregated data can be easily manageable, while individual data might be difficult to manage. A possible solution in this regard would be to analyze only a subset of users, but this would raise the problem of selecting a good and representative sample.

4. TECHNIQUES FOR MOBILE PHONE NETWORK DATA ANALYSIS

In this section, we will show several techniques for mobile phone network data analysis that have been used in research works (some of them are briefly introduced in de Jonge et al. [2012]). First, we will describe some filtering techniques necessary to reduce noise in the data. Then, we will describe a list of features that can be extracted from mobile phone network data as well as the necessary processing techniques.

4.1. Filtering Techniques

In order to mine mobile phone network data to derive human patterns in cities, several techniques are needed to reduce both the spatial uncertainty and the noisiness of the raw data. The main issues in this regard are (1) assigning the user to a specific location and (2) identifying when the user stops in a location or is simply passing through it.

—**Assigning the user to a specific location.** State-of-the-art works in the area suggest two main solutions:
  (1) *Assign the user to the centroid of the cell area*. As shown in Section 3, each CDR produced by a mobile phone is associated to a cell whose location is known by the mobile phone operator. González et al. [2008] first divide the area under investigation with a Voronoi tessellation technique based on the cell tower locations, and then they assign the user position to the centroid of the corresponding Voronoi cell. A different approach is shown in Girardin et al. [2009], where the operator provided the best serving cell map, which associates to each location in a grid the cell tower that best covers that location. The computation is made on simulated coverage and takes into account both the cell sector and propagation models.
  (2) *Assign the user a probability to be in a given location*. This second solution introduces uncertainty in assigning a user to a location. For example, Traag et al. [2011] use a propagation model to assign a user a probability of being at a specific location, given the fact that he or she is connected to a particular cell tower. The main advantage is that this solution takes into consideration the fact that multiple towers might be covering the same location.

—**Stop detection.** Another important issue is determining which places are important to the user, that is, in which places the user stops for a reasonable time period. Given the rawness of mobile phone network data, the same event can be registered as consecutive events associated to different close-by locations. The solutions proposed so far to improve accuracy in the raw mobile phone network data can be divided into two groups:
  (1) **Solutions that leverage consecutive location data,** where consecutive measurements that are close enough can be collapsed in a unique single measurement. For example, Calabrese et al. [2010] fixed both a spatial S_{th} and a temporal T_{th} threshold in order to detect stops; that is, two consecutive stops $stop_i$ and $stop_j$ can be collapsed in the same stop if $distance(d_{stop_i}, d_{stop_i}) < S_{th}$ and $(t_{stop_i} - t_{stop_j}) > T_{th}$. A similar approach was also used in Jiang et al. [2013].
  (2) **Solutions that leverage historical location data,** where historical location data is used to help understand which places are important for the user. For example, Isaacman et al. [2011] use clustering techniques (in particular the Hartigan's algorithm) on a dataset spanning over 78 days with the aim of identifying which places are important to the users, such as home and work location.

4.2. Processing Techniques

In this section, we summarize different techniques proposed in the literature to process mobile phone network data and extract insights into urban dynamics. These techniques have been categorized based on the aggregation level provided by the datasets.

4.2.1. Individual Data Processing. Mobile phone data at the individual level has been used in several applications:

(1) **Home and work location estimation.** Using CDR with location information, some works [Calabrese et al. 2011; Isaacman et al. 2011] have focused on estimating

the home and work location of the users. The technique used to achieve this goal involves selecting, for each user, a dataset consisting of several days of mobile phone network data. Necessary information in the raw data is (1) the number of times a cell tower was contacted by the user and (2) the length (in terms of time) of stay in a location. Home location is then determined as the most frequented place during evenings (where an evening is characterized by a time internal to be specified) and work location as the most frequented place during weekday mornings/afternoons and excluding the home location and places with a high number of evening events. Data has been validated using U.S. Census population estimates at the census track level. Please note that when applying the technique to different countries, the time intervals to be used to identify evening and morning periods might have to be adjusted based on the working habits of the country, as discussed, for instance, in Berlingerio et al. [2013].

(2) **Mobility estimation and applications.** By connecting the sequence of visited locations for each user and using that as an estimation of mobility, several researchers have proposed applications for mobility study. González et al. [2008] proposed a technique to infer daily trips using the distance between any two different visited locations. Distance between the two most distant visited locations has been used in Isaacman et al. [2011] as a measure of daily range of mobility. By grouping users' mobility by origin and destination of trips, Origin–Destination matrices can be inferred and used to analyze the attractiveness of an area (measured as the number of different places people come from); see, for example, Calabrese et al. [2011]. Couronne et al. [2011] cluster users on the basis of how often they move using spatiotemporal analysis. Schneider et al. [2013] associated daily mobility networks extracted from the sequence of trips in a day, called *motifs*, with trip chains extracted from travel diary surveys and tried to associate a trip purpose by examining semantic-enriched land users surrounding destinations of individuals' motifs [Jiang et al. 2013].

Berlingerio et al. [2013] have further exploited frequent travel patterns found in the mobile phone data in order to come up with recommendations for improving public transportation systems by recommending the introduction of new routes in areas that experience high travel demand that is, however, not matched by the current transit network.

Finally, using mobility patterns extracted from CDRs, researchers have defined a model that describes how diseases could spread across the country [Lima et al. 2013]. This has led to testing the effect of information campaigns in containing the disease spread.

(3) **Integrating social and mobility information.** A first group of works tries to understand the interplay of mobility patterns and social ties [Cho et al. 2011; Crandall et al. 2010; Pan et al. 2013; Wang et al. 2011]. As shown in Section 2, mobile phone network data has also been mined to integrate the calling and location pattern in order to help in inferring face-to-face meetings. Calabrese et al. [2011] and Wu et al. [2008] discovered that people calling while connected to the same cell tower (colocation) are a good proxy for face-to-face meetings. In particular, they discovered that people tend to interact much more just before and after this event, and the number of inferred face-to-face meetings decreases with the users' home distance. From the call interactions, the authors are able to predict when and where people will be meeting.

4.2.2. Aggregated Data Processing. As shown in Section 3 compared to individual data, aggregated data is much more easy to manage and can be possibly available in real time. In the following we will show the techniques that have been applied to mobile phone network data in the state-of-the-art works.

(1) **Land use inference.** Starting from aggregated cell tower statistics, it is possible to understand activities in the city from telecommunication usage patterns. This can augment existing built environment data collection and analysis methods (census, business registrations, etc.) at low cost and with very low latencies. Several categories of activities can be considered. Classical time series analysis is initially performed (e.g., the Principal Component Analysis technique has been used in Reades et al. [2007] or the Dynamic Time Warping technique in Yuan and Raubal [2012]) and clustering of time series can classify places based on usage (like the Fuzzy C-Means technique proposed in Soto and Frias-Martinez [2011]).
(2) **Space partitioning.** Mobile phone users' location at call time can be used to infer the location of callers, thus allowing one to model the effect of geography on human interactions. Using network analysis, Lambiotte et al. [2008] found that human interactions decrease as distance increases following a gravity-like behavior. Exceptions emerge and are mainly due to geographical features (e.g., rivers; see, for example, Ratti et al. [2010]), administrative borders, and cultural differences. Using aggregated CDR with location information, one can measure the level of human interactions between places. This has led to several works focused on how to best cluster areas based on these interactions. At the city scale, interaction events can be aggregated to create a network of places where nodes are locations (e.g., cell towers) and edges between nodes exist if interactions happen between people connected to the two cell towers. The weighted graph can be partitioned in communities using standard network analysis techniques (e.g., modularity optimization). Through that, researches can detect which areas in the city are most connected and where interaction borders exist; see Blondel et al. [2010]. An important aspect to take into account while performing this study is the mobile phone penetration and share of the operator in each area under analysis. Indeed, if such share is not uniformly distributed over the entire area under analysis, the resulting interactions network might be distorted. This was one of the problems addressed in Calabrese et al. [2011] when dealing with regional partitioning at the level of the entire United States. Starting from CDR data with location information aggregated at the county level, the authors had to take some actions: (1) normalization in order to deal with operator share not being equal for every area and (2) filtering of counties with a too low number of customers or share (to preserve representativeness of the sample). More recently, new methods have also been proposed to estimate the significance of the association between geographical divisions of the population originating in ethical, language, religious, or political differences [Bucicovschi et al. 2013]. A study has been conducted on the Ivory Coast to take into account the 60 local major languages spoken.
(3) **Event detection.** Looking at time series of call-tower-to-tower communications, researchers have proposed a visual analytic tool to characterize events [van den Elzen et al. 2013]. This tool identifies clusters of cell towers having similar call behavior to detect events. The characterization can be further refined by introducing individual data to identify whether mobile phone users are unusually found in the specific location where the mobile traffic anomaly was detected [Traag et al. 2011]. A threshold to detect these outliers has to be imposed and tuned based on partial ground truth on historical events.

Based on what has been discussed, Table V summarizes our recommendations on which datasets and processing techniques should be used to develop specific urban sensing applications.

Table V. Urban Sensing Applications and Associated Datasets and Processing Techniques

Application	Preferred Dataset	Processing Techniques	Observations
Estimating population density	Individual CDR with cell tower location information	Home location determination	Test different temporal thresholds for home location determination
Estimating types of activities in different parts of the city (e.g., land use)	Aggregated cell tower statistics	Time series clustering (e.g., Fuzzy K-means)	Can be improved with the help of external data, e.g., POIs
Estimating mobility patterns (Origin–Destination matrices)	Individual CDR with cell tower location information	Home and work location determination, mobility estimation	Evaluate feasibility to map match on transport network
Estimating mobility patterns (traffic monitoring)	Individual event-driven triangulated location	Mobility estimation, mode inference	Evaluate availability of data in real time
Detecting events	Individual or aggregated CDR with cell tower location information	Mobility estimation, event identification	Test detection thresholds on partial ground truth
Analyzing the geography of social networks (regional partitioning)	Aggregated CDR with cell tower location information	Modularity partitioning	Test different definitions of weights on edges
Analyzing the geography of social networks (communication–mobility interplay)	Individual CDR with cell tower location information	Mobility estimation, social network analysis	Use reciprocate calls to identify social ties. Use location at call time to identify colocation

5. OPEN CHALLENGES

In this article, we have shown how mobile phone network data can be used to gain insights into urban dynamics. In dealing with this type of data, some challenges still remain:

(1) **Limitations of event-driven data.** In order to analyze certain types of urban patterns, it is important to have very frequent location data. As explained in Section 3, event-driven data are generated only when the user takes some action, for example, sends an SMS, makes a call, and so forth. Thus, the location of the user might not be updated very frequently. Some approaches proposed so far to solve this problem are as follows:
 —*Sampling only highly active users.* This solution might be effective since high communication (e.g., calling someone or sending an SMS) has been found to be correlated to high mobility [Couronne et al. 2011]. The main problem in this regard is how to choose users that represent a good sample of citizens' behavior.
 —*Sampling Internet usage data.* Given the high penetration of smartphones [Manyika et al. 2011], another option is to use the Internet usage data to derive location data. The main pros is that such kind of data generally presents a lower interevent time [Calabrese et al. 2010]; however, smartphone users' behavior does not always represent a general sample of citizens' behavior.
 —*Network-driven data.* Given the low frequency of users' localization updates, a better type of data could be network-driven data. In particular, periodic sampling is independent of events but is not too good for short-term mobility. Another

alternative could be to use mobility-based sampling that is good for analyzing mobility between large areas such as location areas.

(2) **Limitations in spatial accuracy.** It might be important to have very precise location data for certain types of applications, such as to determine the accurate location, the route undertaken by the user, or the transportation modes. As shown in Section 3, mobile phone network data does not provide accurate localization. Some solutions proposed so far are as follows:
—*Look at history for recurring locations.* This can help in smoothing irregularities in the location data, allowing one to assign the nearest recurring location to a noisy position (because of the low accuracy in the localization); see Isaacman et al. [2011].
—*Look at handover during calls.* Handoff patterns are relatively stable across different routes, speeds, directions, phone models, and weather conditions [Becker et al. 2011], thus allowing one to derive the trajectories of mobile devices also using CDR data with a low frequency of localization updates.

(3) **Managing uncertainties.** Looking at the previous open challenges, it is clear that the uncertainties in the user's status in time and space can be relatively large. This is due to both the low frequency of users' localization updates and the spatial resolution of mobile phone network data. Thus, it is important to provide reliable and uncertain-aware results. One proposed solution is estimating uncertainties in users' position. For example, Couronne et al. [2011] try to estimate the bias of user behavior in mobile phone data, taking into account the imprecision of data, with a trigonometric approach to describe both mobility values and uncertainty.

(4) **Finding comparative datasets.** Traditional city data (e.g., census and surveys) are collected using different methods, sampling times, and collection years. This makes it difficult to compare results obtained analyzing mobile phone network data with these traditional datasets. Proposed alternatives are as follows:
—*Self-reported data.* Self-reported data can provide additional value compared to traditional data since they might be more spatially accurate, not outdated, and with a frequent sampling time to make comparisons. An example of self-reported data is that which can be obtained from Flickr,[23] which is used, for example, in Girardin et al. [2008] to mine tourists' patterns in Rome.
—*Social networking data.* Similar to the previous one, social networking data provides specific information regarding the places visited by the users. There are a plethora of location-based social networks such as Foursquare[24] that provide public access to their own data and have been recently used to support urban analysis; see, for instance, Noulas et al. [2013].

(5) **Dealing with privacy and anonymity.** The sharing of mobility data raises serious privacy concerns. Mobility data can reveal the mobility behavior of the people: where they are going, where they live, where they work, their religion preferences, and so forth. All this information refers to the private personal sphere of a person and so may potentially reveal many facets of his or her private life. As a consequence, this kind of data has to be considered personal information to be protected against undesirable and unlawful disclosure. A recent study by de Montjoye et al. [2013] showed that knowing four spatiotemporal points is enough to uniquely identify 95% of the individuals. Thus, sophisticated techniques should be designed to protect the privacy of individuals. Many privacy-enhancing technologies for

[23]http://www.flickr.com.
[24]http://foursquare.com.

mobility data have been proposed by the scientific community; see Giannotti and Pedreschi [2008] for a review on privacy in mobility data. In particular, two proposed solutions [Krumm 2009] so far are as follows:

—*Location obfuscation*, which consists of nonreversible ways to slightly alter the location such that it does not reflect the real location of the user but still contains enough information to provide a satisfactory service. See Wightman et al. [2011] for more information regarding the evaluation of several location obfuscation techniques.

—*k-Anonymity for trajectories*, which ensures that each individual trajectory can only be released if there are at least $k - 1$ distinct individuals whose associated trajectories are indistinguishable from the former (see Gedik and Liu [2008] for more detailed information).

Very recently, Mir et al. [2013] also proposed a method, validated against billions of location samples from a real telecommunication network, to generate synthetic CDRs to capture the mobility patterns of real metropolitan populations while preserving privacy.

This is just the tip of the iceberg. The concerns that people have over the collection of this data will naturally extend to any analytic capabilities applied to the data, even the ones that try to preserve users' privacy. Users of data mining should start thinking about how their use of this technology will be impacted by legal issues related to privacy. A critical evaluation of data mining and privacy was released in a report saying that data mining "may be the most fundamental challenge that privacy advocates will face in the next decade..." [Cavoukian 1998]. The report looks at data mining and privacy in the context of the international "fair information practice" principles.

These collisions between data mining and privacy are just beginning. Over the next few years, we should expect to see an increased level of scrutiny of data mining in terms of its impact on privacy. The sheer amount of data that is collected about individuals, coupled with powerful new technologies such as data mining, will generate a great deal of concern by consumers. Unless this concern is effectively addressed, we expect to see legal challenges to the use of data mining technologies.

(6) **Mobility/communication interplay.** Studying the interplay between telecommunications and physical location is still a challenge. In some cases, it has been suggested that telecommunications may be a substitute for physical interaction [Albertson 1977]. In other cases, conflicting hypotheses have been made, including those of a complementary [Mok et al. 2010], neutral [Choo et al. 2010], or reinforcing effects [Sasakia and Nishiib 2010]. Regarding mobile phone network data, Calabrese et al. [2011] investigate the relationship between people's calls and their physical location. Wang et al. [2011] mine the similarities between people's movements (as collected by the mobile phone network) and social networks. Still, a lot of work has to be done in this area to fully characterize the real interplay.

(7) **Real-time data acquisition and processing.** Many urban sensing applications (e.g., traffic monitoring, event management) are useful if results are presented in real time or near real time. The problem is that usually mobile phone network data is first acquired and then pushed to databases, and thus it is not usually available in real time (see Section 3). Since the quantity of mobile phone network data produced every day is massive, there is a need for ad hoc algorithms and platforms to process such data in real time. Some proposed solutions are based on streaming platforms able to deal with different types of data in real time; see, for example, Kaiser and Pozdnoukhov [2013].

6. CONCLUSIONS

This article discusses the current state of the art and open challenges in the emerging field of mobile phone network data for urban sensing. Telecom operators are nowadays generating terabytes of records of potential use for urban sensing. Research is still particularly needed in (1) inferring behavioral patterns, (2) building analytics and systems to process massive datasets and automatically extract patterns, and (3) building control systems able to make use of inferred patterns to optimize city services. Privacy is also a very sensitive issue that had to be addressed. Mobile phone network data will ultimately provide both micro- and macroscopic views of cities and help understand citizens' behaviors and patterns.

REFERENCES

R. Ahas, S. Silm, O. Järv, E. Saluveer, and M. Tiru. 2010. Using mobile positioning data to model locations meaningful to users of mobile phones. *Journal of Urban Technology* 17, 1 (2010), 3–27.

L. A. Albertson. 1977. Telecommunications as a travel substitute: Some psychological, organizational, and social aspects. *Journal of Communication* 27, 2 (1977), 32–43. DOI:http://dx.doi.org/10.1111/j.1460-2466.1977.tb01824.x

J. P. Bagrow, D. Wang, and A.-L. Barabási. 2011. Collective response of human populations to large-scale emergencies. *PLoS ONE* 6, 3 (2011).

R. Becker, R. Cáceres, K. Hanson, S. Isaacman, J. M. Loh, M. Martonosi, J. Rowland, S. Urbanek, A. Varshavsky, and C. Volinsky. 2013. Human mobility characterization from cellular network data. *Communications of the ACM* 56, 1 (Jan. 2013), 74–82. DOI:http://dx.doi.org/10.1145/2398356.2398375

R. A. Becker, R. Caceres, K. Hanson, J. M. Loh, S. Urbanek, A. Varshavsky, and C. Volinsky. 2011. Route classification using cellular handoff patterns. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp'11)*. ACM, New York, NY, 123–132. DOI:http://dx.doi.org/10.1145/2030112.2030130

M. Berlingerio, F. Calabrese, G. Lorenzo, R. Nair, F. Pinelli, and Marco L. Sbodio. 2013. AllAboard: A system for exploring urban mobility and optimizing public transport using cellphone data. In *Machine Learning and Knowledge Discovery in Databases*, Hendrik Blockeel, Kristian Kersting, Siegfried Nijssen, and Filip Železný (Eds.). Lecture Notes in Computer Science, Vol. 8190. Springer, Berlin, 663–666. DOI:http://dx.doi.org/10.1007/978-3-642-40994-3_50

V. Blondel, G. Krings, and I. Thomas. 2010. Regions and borders of mobile telephony in Belgium and around Brussels. *Brussel Studies* 42 (2010).

O. Bucicovschi, R. W. Douglass, D. A. Meyer, M. Ram, D. Rideout, and D. Song. 2013. Analyzing social divisions using cell phone data. In *Proceedings of the 3rd International Conference on the Analysis of Mobile Phone Datasets (NetMob'13)*. Boston, MA.

F. Calabrese, M. Colonna, P. Lovisolo, D. Parata, and C. Ratti. 2011a. Real-time urban monitoring using cell phones: A case study in rome. *IEEE Transactions on Intelligence Transportation Systems* 12, 1 (2011), 141–151.

F. Calabrese, D. Dahlem, A. Gerber, D. Paul, C. Xiaoji, J. Rowland, C. Rath, and C. Ratti. 2011b. The connected states of America: Quantifying social radii of influence. In *Privacy, Security, Risk and Trust (SocialCom'11)*.

F. Calabrese, G. Di Lorenzo, Liang Liu, and C. Ratti. 2011c. Estimating origin-destination flows using mobile phone location data. *IEEE Pervasive Computing* 10, 4 (April 2011), 36–44. DOI:http://dx.doi.org/10.1109/MPRV.2011.41

F. Calabrese, F. Pereira, G. Di Lorenzo, L. Liu, and C. Ratti. 2010. The geography of taste: Analyzing cellphone mobility and social events. In *Pervasive Computing*, P. Floréen, A. Krüger, and M. Spasojevic (Eds.). Lecture Notes in Computer Science, Vol. 6030. Springer, Berlin, 22–37. DOI:http://dx.doi.org/10.1007/978-3-642-12654-3_2

F. Calabrese, Z. Smoreda, V. D. Blondel, and C. Ratti. 2011. Interplay between telecommunications and face-to-face interactions: A study using mobile phone data. *PLoS ONE* 6, 7 (07 2011), e20814. DOI:http://dx.doi.org/10.1371/journal.pone.0020814

A. Cavoukian (Ed.). 1998. *Data Mining: Staking a Claim on Your Privacy*. Information and Privacy Commissioner/Ontario.

E. Cho, S. A. Myers, and J. Leskovec. 2011. Friendship and mobility: User movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge*

Discovery and Data Mining (KDD'11). ACM, New York, NY, 1082–1090. DOI:http://dx.doi.org/10.1145/2020408.2020579

S. Choo, T. Lee, and P. L. Mokhtarian. 2010. Do transportation and communications tend to be substitutes, complements, or neither?: U.S. consumer expenditures perspective, 1984-2002. *Transportation Research Record: Journal of the Transportation Research Board* 2010/2007, 1 (2010), 121–132.

T. Couronne, A.-M. Olteanu, and Z. Smoreda. 2011a. Urban mobility: Velocity and uncertainty in mobile phone data. In *SocialCom/PASSAT'11*. 1425–1430.

T. Couronne, Z. Smoreda, and A.-M. Olteanu. 2011b. Chatty mobiles: Individual mobility and communication patterns. *NetMob 2011* (2011).

D. J. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg. 2010. Inferring social ties from geographic coincidences. *Proceedings of the National Academy of Sciences* 107, 52 (Dec. 8, 2010), 22436–22441. DOI:http://dx.doi.org/10.1073/pnas.1006155107

E. de Jonge, M. van Pelt, and M. Roos. 2012. Time patterns, geospatial clustering and mobility statistics based on mobile phone network data. *Statistics Netherlands, Division Methodology and Quality Sector Methodology Heerlen* (2012).

Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports* 3 (March 25, 2013). DOI:http://dx.doi.org/10.1038/srep01376

L. Ferrari, M. Mamei, and M. Colonna. 2012. People get together on special events: Discovering happenings in the city via cell network analysis. In *PerCom Workshops*. 223–228.

R. L. Freeman. 2005. *Fundamentals of Telecommunications*. John Wiley.

United Nations Population Fund. 2007. 2007 State of World Population. Retrieved from http://www.unfpa.org/swp/swpmain.htm.

B. Gedik and L. Liu. 2008. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing* 7, 1 (Jan. 2008), 1–18. DOI:http://dx.doi.org/10.1109/TMC.2007.1062

F. Giannotti and D. Pedreschi (Eds.). 2008. *Mobility, Data Mining and Privacy: Geographic Knowledge Discovery*. Springer.

F. Girardin, F. Calabrese, F. D. Fiore, C. Ratti, and J. Blat. 2008. Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing* 7, 4 (Oct.-Dec. 2008), 36–43. DOI:http://dx.doi.org/10.1109/MPRV.2008.71

F. Girardin, A. Vaccari, A. Gerber, and C. Ratti. 2009. Quantifying urban attractiveness from the distribution and density of digital footprints. *Journal of Spatial Data Infrastructure Research* 4 (2009), 175–200.

M. C. González, C. A. Hidalgo, and A.-L. Barabási. 2008. Understanding individual human mobility patterns. *Nature* 453, 7196 (June 2008), 779–782. DOI:http://dx.doi.org/10.1038/nature06958

International Telecommunication Union. Homepage. Retrieved from http://www.itu.int/ITU-D/ict/.

S. Isaacman, R. Becker, R. Cáceres, S. Kobourov, M. Martonosi, J. Rowland, and A. Varshavsky. 2011. Identifying important places in people's lives from cellular network data. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive'11)*. Springer-Verlag, Berlin, Heidelberg, 133–151.

S. Isaacman, R. Becker, R. Cáceres, S. Kobourov, J. Rowland, and A. Varshavsky. 2010. A tale of two cities. In *Proceedings of the 11th Workshop on Mobile Computing Systems Applications (HotMobile'10)*. ACM, New York, NY, USA, 19–24. DOI:http://dx.doi.org/10.1145/1734583.1734589

S. Jiang, G. A. Fiore, Y. Yang, J. Ferreira, Jr., E. Frazzoli, and M. C. González. 2013. A review of urban computing for mobile phone traces: Current methods, challenges and opportunities. In *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing (UrbComp'13)*. ACM, New York, NY. DOI:http://dx.doi.org/10.1145/2505821.2505828

C. Kaiser and A. Pozdnoukhov. 2013. Enabling real-time city sensing with kernel stream oracles and mapreduce. *Pervasive and Mobile Computing* 7, 5 (2013), 708–721.

J. M. Krisp. 2010. Planning fire and rescue services by visualizing mobile phone density. *Journal of Urban Technology* 17, 1 (2010), 61–69.

J. Krumm. 2009. A survey of computational location privacy. *Personal Ubiquitous Computing* 13, 6 (Aug. 2009), 391–399. DOI:http://dx.doi.org/10.1007/s00779-008-0212-5

R. Lambiotte, V. D. Blondel, C. de Kerchove, E. Huens, C. Prieur, Z. Smoreda, and P. Van Dororen. 2008. Geographical dispersal of mobile communication networks. *Physica A: Statistical Mechanics and Its Applications* 387 (2008), 5317–5325.

A. Lima, M. De Domenico, V. Pejovic, and M. Musolesi. 2013. Exploiting cellular data for disease containment and information campaigns strategies in country-wide epidemics. In *Proceedings of the 3rd International Conference on the Analysis of Mobile Phone Datasets (NetMob'13)*. Boston, MA, USA.

X. Lu, L. Bengtsson, and P. Holme. 2012. Predictability of population displacement after the 2010 Haiti earthquake. *Proceedings of the National Academy of Sciences* 109, 19 (2012).

J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers. 2011. Big data: The next frontier for innovation, competition and productivity. In *McKinsey Global Institute*.

D. J. Mir, S. Isaacman, R. Caceres, M. Martonosi, and R. N. Wright. 2013. DP-WHERE: Differentially private modeling of human mobility. In *Proceedings of the 2013 IEEE International Conference on Big Data*. 580–588. DOI:http://dx.doi.org/10.1109/BigData.2013.6691626

D. Mok, B. Wellman, and J. Carrasco. 2010. Does distance still matter in the age of the internet? *Urban Studies* 47, 13 (2010), 2747–2783.

J. More and C. Lingam. 2013. Current trends in reality mining. *IRJES* 2, 2 (2013), 35–39.

A. Noulas, C. Mascolo, and E. Frías-Martínez. 2013. Exploiting foursquare and cellular data to infer user activity in urban environments. In *MDM*. 167–176.

P. Nurmi and S. Bhattacharya. 2008. Identifying meaningful places: The non-parametric way. In *Pervasive Computing*, J. Indulska, D. Patterson, T. Rodden, and M. Ott (Eds.). Lecture Notes in Computer Science, Vol. 5013. Springer, Berlin, 111–127.

J.-P. Onnela, J. Saramäki, J. Hyvonen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. 2007. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences* 104, 18 (2007), 7332.

W. Pan, G. Ghoshal, C. Krumme, M. Cebrian, and A. Pentland. 2013. Urban characteristics attributable to density-driven tie formation. *Nature Communications* 4, 1961 (2013).

C. Ratti, S. Sobolevsky, F. Calabrese, C. Andris, J. Reades, M. Martino, R. Claxton, and S. H. Strogatz. 2010. Redrawing the map of great britain from a network of human interactions. *PLoS One* 5, 12 (2010), e14248. DOI:http://dx.doi.org/10.1371/journal.pone.0014248

J. Reades, F. Calabrese, A. Sevtsuk, and C. Ratti. 2007. Cellular census: Explorations in urban data collection. *IEEE Pervasive Computing* 6, 3 (2007), 30–38.

K. Sasakia and K. Nishiib. 2010. Measurement of intention to travel: Considering the effect of telecommunications on trips. *Transportation Research Part C: Emerging Technologies* 18, 1 (2010), 36–44.

C. M. Schneider, V. Belik, T. Couronne, Z. Smoreda, and M. C. Gonzalez. 2013. Unravelling daily human mobility motifs. *Journal of the Royal Society, Interface* 10, 84 (2013).

A. Sevtsuk and C. Ratti. 2010. Does urban mobility have a daily routine? Learning from the aggregate data of mobile networks. *Journal of Urban Technology* 17, 1 (2010), 41–60.

F. Simini, M. C. González, A. Maritan, and A.-L. Barabási. 2012. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96–100.

T. Sohn, A. Varshavsky, A. LaMarca, M. Chen, T. Choudhury, I. Smith, S. Consolvo, J. Hightower, W. Griswold, and E. de Lara. 2006. Mobility detection using everyday GSM traces. In *UbiComp 2006: Ubiquitous Computing*, Paul Dourish and Adrian Friday (Eds.). Lecture Notes in Computer Science, Vol. 4206. Springer, Berlin, 212–224.

C. Song, T. Koren, P. Wang, and A.-L. Barabási. 2010a. Modelling the scaling properties of human mobility. *Nature Physics* 6, 10 (2010), 818–823.

C. Song, Z. Qu, N. Blumm, and A.-L. Barabási. 2010b. Limits of predictability human mobility. *Science* 327, 2 (2010), 1018–1021.

V. Soto and E. Frias-Martinez. 2011. Robust land use characterization of urban landscapes using cell phone data. In *Proceedings of the 1st Workshop on Pervasive Urban Applications (Purba'11)*.

V. Soto, V. Frias-Martinez, J. Virseda, and E. Frias-Martinez. 2011. Prediction of socioeconomic levels using cell phone records. In *Proceedings of the 19th International Conference on User Modeling, Adaption, and Personalization (UMAP'11)*. Springer-Verlag, Berlin, 377–388.

Technology Review. 2010. Mobile Data: A Gold Mine for Telcos. Retrieved from http://www.technologyreview.com/news/419101/mobile-data-a-gold-mine-for-telcos/.

V. A. Traag, A. Browet, F. Calabrese, and F. Morlot. 2011. Social event detection in massive mobile phone data using probabilistic location inference. In *Privacy, Security, Risk and Trust (Passat), 2011 IEEE 3rd International Conference on and 2011 IEEE 3rd International Conference on Social Computing (Socialcom)*. 625–628. DOI:http://dx.doi.org/10.1109/PASSAT/SocialCom.2011.133

S. van den Elzen, J. Blaas, D. Holten, J.-K. Buenen, J. J. van Wijk, R. Spousta, A. Miao, S. Sala, and S. Chan. 2013. Exploration and analysis of massive mobile phone data: A layered visual analytics approach. In *Proceedings of the 3rd International Conference on the Analysis of Mobile Phone Datasets (NetMob'13)*. Boston, MA, USA.

D. Wang, D. Pedreschi, C. Song, F. Giannotti, and A.-L. Barabási. 2011. Human mobility, social ties, and link prediction. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge*

Discovery and Data Mining (KDD'11). ACM, New York, NY, USA, 1100–1108. DOI:http://dx.doi.org/10.1145/2020408.2020581

P. Wang, T. Hunter, A. M. Bayen, K. Schechtner, and M. C. Gonzalez. 2012. Understanding road usage patterns in urban areas. *Scientific Reports* 2 (2012).

S. Wang, J. Min, and B. Yi. 2008. Location based services for mobiles: Technologies and standards. In *IEEE ICC*. Beijing.

A. Wesolowski, N. Eagle, A. M. Noor, R. W. Snow, and C. O. Buckee. 2013. The impact of biases in mobile phone ownership on estimates of human mobility. *Journal of the Royal Society, Interface* 10, 81 (2013).

P. Wightman, W. Coronell, D. Jabba, M. Jimeno, and M. Labrador. 2011. Evaluation of location obfuscation techniques for privacy in location based information systems. In *2011 IEEE Latin-American Conference on* Communications (LATINCOM). 1–6. DOI:http://dx.doi.org/10.1109/LatinCOM.2011.6107399

L. Wu, B. N. Waber, S. Aral, E. Brynjolfsson, and A. Pentland. 2008. Mining face-to-face interaction networks using sociometric badges: Predicting productivity in an IT configuration task. In *ICIS*. Association for Information Systems, 127.

Y. Yuan and M. Raubal. 2012. Extracting dynamic urban mobility patterns from mobile phone data. In *Proceedings of the 7th International Conference on Geographic Information Science (GIScience'12)*.

Received April 2013; revised March 2014; accepted July 2014