

# Chapter 3

## Observing Human Activity Through Sensing

Sidharta Gautama, Martin Atzmueller, Vasillis Kostakos, Dominique Gillis, and Simo Hosio

### 3.1 Introduction

In this chapter we discuss new ways to observe people's activity, i.e. people's movements and whereabouts. These observations can be valuable for studying mobility, people's movement patterns through the urban environment, their use of the urban space, and finally social interaction.

Following Moore's Law, the progression in processing power is introducing technology into society at an increasing pace. Whereas smartphones at their introduction 15 years ago carried MHz processors with a few Mb of memory, current smartphones are equipped with high-end multicore processors, rich storage, multiple sensors, touch screen and different networking capabilities. The power of these devices equals that of low-end personal computers, but at a lower cost and with better portability. Their uptake has been spectacular, filling a need in our digital society. Today this means that through their phone, people continuously carry in their daily life sensors and computing power which is an important step towards ubiquitous computing as thought out at Xerox PARC beginning of the 1990s (Pentland 2000).

---

S. Gautama (✉) • D. Gillis  
Ghent University, i-KNOW, St.Pietersnieuwstraat 41, B-9000, Gent, Belgium  
e-mail: [sidharta.gautama@ugent.be](mailto:sidharta.gautama@ugent.be)

M. Atzmueller  
University of Kassel, Research Center for Information System Design, Wilhelmshöher Allee 73,  
34121, Kassel, Germany  
e-mail: [atzmueller@cs.uni-kassel.de](mailto:atzmueller@cs.uni-kassel.de)

V. Kostakos • S. Hosio  
University of Oulu, Center for Ubiquitous Computing, Erkki Koiso-Kanttilan katu 3, door E,  
P.O Box 4500, FI-90014, Oulu, Finland  
e-mail: [vassilis@ee.oulu.fi](mailto:vassilis@ee.oulu.fi)

At the other end of the spectrum, low-cost, low-power and miniaturized processors are stimulating the widespread use of sensor nodes. These sensor nodes typically consist out of sensing, processing and communication modules. They can work stand-alone but their added-value comes when building networks of sensor nodes that collaborate on communication and performing tasks. As nodes do not necessarily belong to closed networks but can form a heterogeneous system of interconnected networks, this forms what is called the Internet of Things (IoT) (Atzori et al. 2010). This refers to connected devices that can observe, understand and act upon certain events without human intervention.

These technological evolutions today offer new ways for observing human behavior. In this chapter, we will look at three technologies for observation: scanning, location-enabled devices and tagging, and illustrate with examples how they can be used in travel behavior studies, in characterization of urban space or in the study of social interaction.

## 3.2 Observation by Scanning

Broadly speaking, by scanning technologies we refer to data measured by means of detectors that scan specific locations of interest. They are typically used for traffic monitoring where they are located along the roadside and deployed to capture roadside motorized travel behavior. Generally, these technologies can be split into two categories: intrusive and non-intrusive methods. The intrusive methods basically consist of a data recorder and a sensor placed on or under the surface. They have been employed for many years in a traffic context and the most important ones are pneumatic road tubes, piezoelectric sensors and magnetic loops. This has been widely deployed over the last decades but the implementation and maintenance costs can be expensive.

Non-intrusive methods are based on remote observations. Manual counting is the most traditional method, where trained observers gather traffic data that cannot be efficiently obtained through automated counts e.g. vehicle occupancy rate, pedestrians and vehicle classifications. In addition, other techniques have emerged based on sensing modalities like radar, infrared, ultrasound and video.

The above detectors are currently in operation for traffic count operations, where the focus lies on estimating traffic volume, possibly annotated with speed and vehicle class and aimed towards motorized transport. Mobility studies today require however a higher level of detail, giving views on (1) network connections and travel flows and (2) all users of the mobility network including pedestrians, bicyclists, public transport users. This requires not only observation of an object in a single location. It becomes necessary to follow an object over a network or site, if not completely than at least for sampled observations of its path. In this section we describe camera networks and bluetooth scanning as examples of this point-to-point scanning technology.

### 3.2.1 *Computer Vision and Camera Networks*

Closed-circuit television (CCTV) systems have known a widespread use since the 1970s with its main application in surveillance and security. During this time, the technology of video cameras and recording has known significant advances due to the evolution of sensors, computing power and digital transmission. This led among others to the migration from analogue cameras to digital technology, the utilization of the Internet Protocol (IP) for video and remote monitoring and the increased use of pan-tilt-zoom control (Kruegle 2011). On the video analytics side, better algorithms and increased computing power on the server as well as embedded on the camera has led to more intelligent applications in 2D and 3D.

Highway traffic camera systems for speeding and toll charging are based on automatic number plate reading (ANPR). These systems consist out of three processes: image capture, plate extraction and interpretation. The performance is reported to have a 90–94% overall read accuracy under optimal conditions. Errors can be introduced due to bad calibration, bad lighting conditions, obscurity and processing ambiguity. Vehicle classification has been extensively studied for highway and urban traffic settings as reviewed by Buch et al. (2011). Most of the related works deal with the vehicle classification problem under good and steady illumination conditions. More challenging scenarios for urban settings have been studied by comparing the vehicle silhouettes against projected 3D models of several vehicle classes. Edge maps, SIFT descriptors, and region-based features are the most common methods employed in the literature to describe the vehicle appearance. In industrial traffic monitoring solutions like Honeywell and Flir, we also find vehicle classification which is often performed under more controlled conditions (e.g. camera view, observed driving) in order to simplify the processing for speed and robustness. Figure 3.1 illustrates such a system.

For machines to be able to detect, track and identify people instead of vehicles is more challenging. As people behave in a more erratic way and have much more variation in appearance, sensing of humans has long been one of the hardest machine vision problems to tackle. Success came from a combination of well-established pattern recognition techniques with an understanding of the image generation process. These methods often capitalized on regularities that are peculiar to people, as for instance human skin color which is defined on a one-dimensional manifold or the human facial geometry.

There are now several companies that sell commercial face recognition software that is capable of high-accuracy recognition with a database of over 1000 people, commercially available camera systems that perform real-time face tracking for teleconferencing, and companies like IBM, Microsoft, Mitsubishi and Sony are showing simple vision-based recognition interfaces in commercial applications. Bredereck et al. (2012) performs tracking of people in camera networks by first detecting persons using histograms of oriented gradients in each camera view and then tracking their positions with a particle filter and greedy matching. In (Morbee et al. 2010), the concept of probabilistic occupancy mapping is utilized to locate



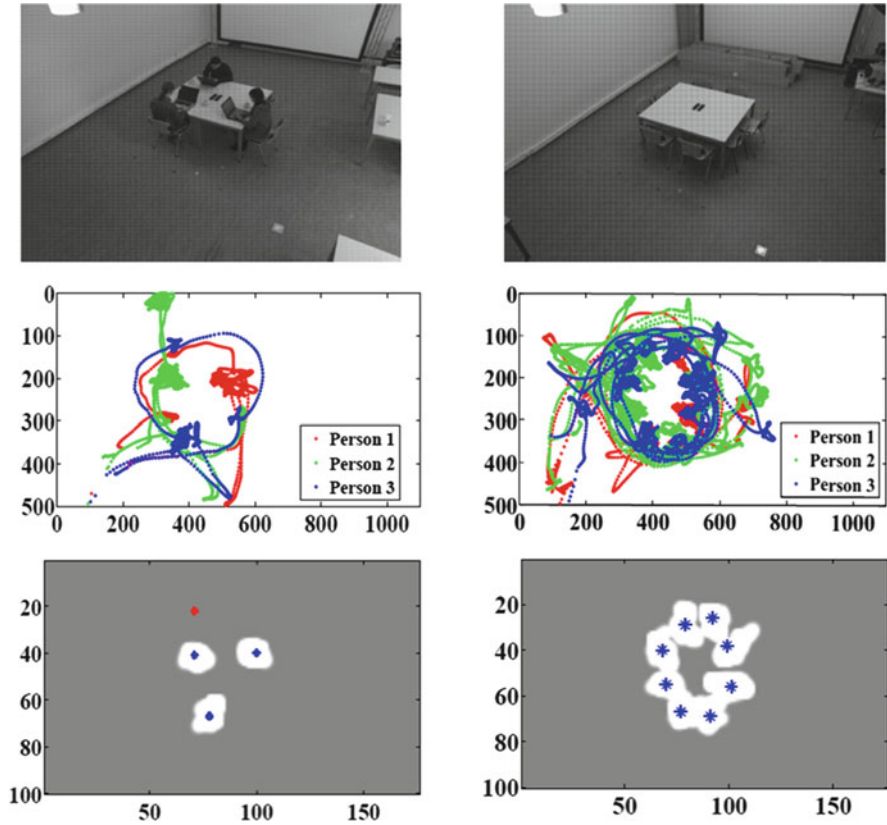
**Fig. 3.1** Example of system for vehicle counting and classification (courtesy of FLIR)

persons in each view and track/fuse using optimization techniques. Figure 3.2 shows results from Xie et al. (2012), which studies the classification of activities into three categories (sitting, standing and walking) based on the estimated trajectories of people in order to infer the position of furniture in a room. The above techniques allows face- or silhouette-based identification and tracking of people from camera to camera which scale up to the monitoring of rooms or small buildings. On larger scales (e.g. urban areas where 10,000+ people pass daily), performance will drop due to limits in recognition.

### 3.2.2 *Bluetooth Scanning*

More recently, Bluetooth has been suggested as an interesting alternative tracking technology. Since the Bluetooth protocol allows for wireless discovery and identification of nearby devices, static Bluetooth sensors placed at strategic locations can give insights into human mobility in a variety of contexts: dynamics at mass events, urban design, social studies, travel time estimation of motorized traffic (Versichele et al. 2012; Eagle and Pentland 2006).

Initially envisioned as a low-power and open protocol for implementing Wireless Personal Area Networks by Siemens in 1994, Bluetooth has since become an almost ubiquitous technology on modern mobile devices. Prior to the ability for two devices to connect wirelessly through Bluetooth, one device needs to be discovered by the



**Fig. 3.2** Learning activities (sitting, walking) from observed trajectories of people and inferring the position of chairs (Xie et al. 2012)

other. This part of the Bluetooth protocol is called the inquiry phase. The master device transmits inquiry packets, to which discoverable devices within its vicinity respond with inquiry response packets. These include the MAC address (which is a 48-bit identifier of the mobile device), and the class of device (COD) code (which gives a general idea about the type of device and some of its functionalities). By mapping detected MAC addresses to a specific timestamp and location where a sensor that made the discovery was located, one can reconstruct proximity-based trajectories. Since an actual connection is not required, tracked individuals are not aware of the presence of Bluetooth sensors and the methodology is in essence completely unobtrusive. Since Bluetooth 1.2, it is also possible to register the received signal strength indicator (RSSI) of the inquiry response packets, which is loosely correlated with the distance between the sensor and the detected device.

Several studies report on the tracking of Bluetooth sensors over a study area and reconstructing movements by matching the MAC addresses of detected devices with the locations of the detecting sensors. This network-based and non-participatory

approach started materializing after a first documented trial in (O'Neill et al. 2006). Since then, a growing number of experimental use-cases have been documented. Particular attention has been devoted to the use of Bluetooth technology for travel time measurements of motorized traffic as it represents a simplified approach in comparison to ANPR (Malinovskiy et al. 2010). Pedestrian mobility has also been investigated. Examples include transit time measurements in airport security checkpoints (Bullock et al. 2010), travel and dwelling time calculations in an urban context (Malinovskiy et al. 2010), and the automatic registration of public transport users (Weinzerl and Hagemann 2007). The 'Cityware' project used static Bluetooth sensors to capture mobility traces, and coupled these data with user's online social data (Kostakos et al. 2008).

An example of Bluetooth scanning for monitoring crowd behavior during large scale city events is described in (Versichele et al. 2012). Bluetooth scanning has been implemented to the Ghent Festivities event in Ghent (Belgium) in support of city event management for the organization, security, transport and emergency service providers. This event takes place on 11 squares in the city center, acting as major attractors (on-stage performances, bars and food stands, fairs). During the festivities 22 locations were covered with Bluetooth scanners, representing people's mobility within the festivity zone itself (trips from one square to another), but also the mobility to and from the festivity zone (trips to and from the two main train stations in Ghent and one park and ride facility).

Applications of the resulting data are manifold. The most direct result are the statistics about visitors and their behavior, such the number of visitors per day, the total number of visitors and unique visitors, the distribution of visitors during the day, the (sequence of) square visited by individual visitors, etc. A second, derived set of results deals with the distribution of the crowd in the festivity zone and the flows of visitors in the city center. This information is vital for security services, which are monitoring the people density in order to plan safety measures as temporary closures of access to overcrowded squares or facilitating the circulation between certain squares. This information is made available to visitors by the festivity app, assisting them to plan their day avoiding crowded routes and squares. A third set of results deals with the trip to and from the festivity zone. A smart choice of detector locations, gives insight in the modal split (train users at the train stations, car users at the park and ride facility, tram users at tram stops). Monitoring the travel times between these locations and the festival zone indicates the accessibility of the event. For example high dwelling times at the park and ride facility point to long waiting time to get onto the tram towards the city center, urging the public transport company to (temporarily) increase the capacity on this line. High travel times between the park and ride facility and the city center suggest congestion problems on the route, where traffic police should intervene to facilitate the circulation of the tram. In this way the Bluetooth scanners assist partners to optimize safety and comfort to the visitors of the festivity. In terms of data quality, care should be taken that Bluetooth scanning samples the population through the observed activity of discoverable Bluetooth devices. If absolute density or flow statistics are needed, this requires extrapolation by estimating the ratio BT versus non-BT carrying visitors at

the scanner positions. These ratios can be estimated by counting the total number of visitors, either manually or automatically if possible (e.g. vehicle counting using ANPR or traffic counting loops). These are expensive processes, however the more accurate the ratios are estimated, the better the data quality.

### 3.3 Movement Tracking and Presence Detection Using Location-Enabled Devices

Recent technological developments have produced a range of digital tracking technologies that offer a view on the movement of users. Various technologies have been used to capture urban movement across a city, including geotagged photos, mobile phone logs, smart card records, taxi/bus GPS traces, and Bluetooth sensing (Calabrese et al. 2011a; Girardin et al. 2008; Kostakos et al. 2010; Quercia et al. 2011). Movement patterns can be used in mobility studies. It is also possible to develop a better understanding of city-dwellers' space use over time, and subsequently inform important decisions about development, growth, and investment across a city. Understanding how various groups of people move in a particular area, and when, provides better context for understanding the types of potential audiences for services in those areas, but also in terms of long-term investment and development decisions (Quercia et al. 2011).

Location accuracy and power consumption have improved by better signal processing techniques and tighter integration with various technologies (GSM, GNSS, WiFi, motion sensors). Tracking technology is currently tightly integrated with current mobile phones and personal navigation devices (PND) offering various location based services. The simple and standard solution is GNSS-based devices, carrying chipsets that receive and correlate incoming satellite signals from GPS, GLONASS and Galileo for positioning. When incorporating this data in studies on spatial human behaviour, care should be taken in the pre-processing of these data streams. Significant errors will be present due to different possible error sources: (1) poor satellite visibility or reflections (e.g. in urban canyons) leading to positional errors; (2) start up time of the GPS chip on cold start leading to several minutes of missing data at the start of the journey; (3) missing data due to memory overload or communication failure. These errors are propagated when measurements (e.g. speed, acceleration) or higher level information (e.g. trip segments) are estimated from this data. The GPS quality flags are not always a good indicator for the occurrence of these errors and care should be taken when processing the data, either by preprocessing the data (e.g. kalman filtering for data integration) or by calculating more reliable quality measures on the data.

Less known is the fact that cell communication offers other possibilities to track people continuously. Operating on a phone network requires the network operator to be able to detect the subscriber's proximity to a specific antenna, even when no calls are made. Using multilateration of radio signals between antenna masts on



the cellular network and the cell phone, its location can be estimated. In general, the accuracy of tracked mobile devices is lower than GNSS-based devices, ranging from 50 to 100 m depending on the density of the cellular network. Projects like MIT's Senseable City have investigated behavior patterns through cell phone activity (Calabrese et al. 2011a). The analyzed activity is still limited to presence detection within cells of typically 100 m radius and does not take into account dynamic spatial movement patterns. Coarse-grained mobile phone traces are primarily suited for developing models for traffic estimation in some cases in real-time (Calabrese et al. 2011b), and to estimate the precision, metering frequency and the number of localizations necessary to achieve accurate traffic descriptions. The literature has demonstrated clear patterns emerging on a daily and weekly basis across a city.

In a similar way, the WiFi network can be used to locate WiFi-enabled devices within a position calibrated network as is currently offered as a service by Google in Android.

### ***3.3.1 Monitoring Traffic with Floating Car Data***

Floating car data (FCD) come from so-called probe-vehicles, i.e. vehicles that are equipped with the necessary devices to transmit data to a data center at regular time intervals. The data comprise information on the status of the vehicle, for instance its location and speed. The equipment in the probe-vehicles is typically GSM communication sending out a GNSS positioning signal. This can be a simple black box datalogger, as used in fleet management systems, or can be integrated within an internal or external navigation system in the car. In the data center, data is processed in order to make it useable. The accuracy of the derived information depends on the frequency of the positioning and broadcasting of the data, the accuracy of the GPS and the number of probe-vehicles.

The utilization of FCD has been extensively investigated in a number of papers (Asmundsdottir et al. 2010). In The Netherlands, the Ministry of Transport, Public Works and Water Management investigated the usefulness of FCD to get an understanding of the possibilities and problems. The experiment was part of a large innovation research program called "Roads to the Future". Approximately 60 vehicles in the city of Rotterdam were equipped with GPS to estimate travel times. Results indicate that the FCD system is fairly accurate and can be applied to traffic information and traffic management systems (Taale et al. 2000). Torp and Lahrman (2005) propose a complete prototype system that uses FCD for both automatic and manual detection of queues in traffic. The automatic detection was based on analyzing GPS data from the taxis. The manual detection was based on taxi drivers reporting traffic queues by using the equipment in the taxis. Reinthaler et al. (2007) proposed a system that uses FCD to calculate detailed routes and travel times for hazardous goods transport in the Austrian road network.

FCD is used in the production and maintenance of road network databases. This production process requires a lot of work and resources which in addition needs



to be in constant update. Instead of manual surveying or mobile mapping, vehicles can be used as proxy in the sense that where there are vehicles, there must be a road. Companies like TomTom, who have altered their map production process in order to receive and process anonymized GPS-data from the navigation devices of their customers. With this data, they are able to build and update their road network database. In addition to simple, static geometry, they are able to extract information on driving direction, speed limits and dynamic travel time.

### ***3.3.2 Monitoring Dynamic, Multimodal Crowd Behaviour Through Cell Phone Localization and Activity***

The introduction of GPS started an important revolution in travel behaviour studies, i.e. in the collection of trip data. By logging the tracking data, an enormous amount of detailed information on the exact position of devices became available, together with the corresponding time stamps, showing potential to fill some of the gaps that were not covered in traditional methods. However, GPS-logging has a major restriction: as devices are typically installed in vehicles, they only monitor the use of the vehicle and therefore typically cover only a unimodal part of an individual's trip behaviour.

The application of portable handheld GPS-devices offered a solution to this issue, but again required the effort and discipline from the respondent to continuously carry the device with him, as forgetting the device would result in unreported gaps in the trip data. The resulting, multimodal tracking data also introduced the new challenge of interpreting the data. In case of passive logging, where tracking is performed without additional input from the user, the survey does not include any information about trip purpose or travel mode. These characteristics can be reconstructed afterwards, either by means of additional surveys (Asakura and Hato 2004) or by interpreting the data using logical rules, e.g. using speed or GIS information (Tsui and Shalaby 2006). Splitting the continuous GPS-logging into separate trips by detecting origins and destinations is based on dwelling times at one location. The determination of the travel mode is primarily based on speed characteristics during the trip, which can be complemented with additional GIS-data e.g. about public transportation networks and rail networks or accelerometer data (Hato 2010). Trip purpose can be estimated using land use maps or by analysing the individual trip chaining. Good results are achieved for determining trip ends and travel modes, but the estimation of trip purposes remains unsatisfying (Gong et al. 2011).

Smartphones bring new possibilities for tracking, having the same capabilities as the portable GPS-device but with additional sensors which can offer a more solid base for travel mode determination (accelerometer, Bluetooth, WiFi). Carrying a smartphone has also become a habit and is therefore considered less of a burden, reducing the risk of non-reported trips. Furthermore, smartphones offer

the opportunity of running interactive mobile applications, where respondents can report additional trip data, e.g. on trip purpose or travel mode. Although the app requires a manual intervention by the respondent (“active” or “interactive logging”), the burden is limited because the reporting is restricted to short entries at the very moment of departure and arrival. As a consequence, time and location of the departure and arrival can be more accurately detected.

A smartphone application called CONNECT is used in Vlassenroot et al. (2013) to monitor dynamic, multimodal crowd behavior through smartphone localization and activity. In order to collect data for crowd behavior analysis, a mobile application released through the appropriate channel is installed by the user. After registration and informing the user of the purpose of data collection, this application collects different kinds of valuable information and sends it to a central server. If the smartphone contains a GPS-chip, accurate locations of the phone can be collected. However, because GPS is very demanding on the battery and does not work inside buildings, the GPS is only activated when appropriate. Other sensor information can be used to derive the location of the phone as accurately as possible, such as cell towers and WiFi. On top of that, measurements of the accelerometer of the smartphones are collected in order to automatically distinguish between different transport modes (pedestrians, cyclists and cars). Depending on the data campaign, the user can have full control over when and which data is being sent. Because of the diverse positioning modes, tracking can be done indoor as well as outdoor and the battery life can be preserved better than in standard GPS-mode. The CONNECT mobile application is shown in Fig. 3.3.

The application has two modes of operation. The first mode is a passive mode where the application runs as a data logger in the background on the smartphone.

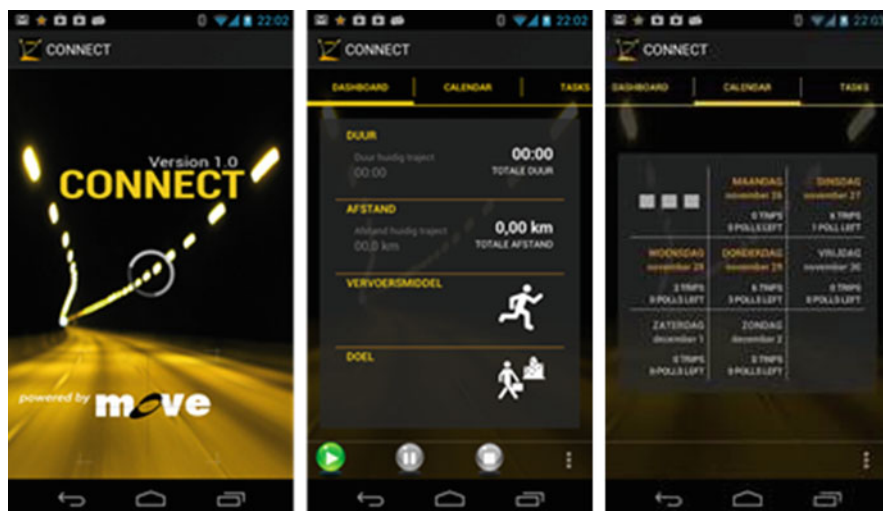



Fig. 3.3 Illustration of the CONNECT mobile application in support of mobility studies

A second mode is a diary mode aimed towards specific mobility studies. For these trials, data acquisition is done by a dedicated test audience carrying a digital diary. The same CONNECT back end is used but with a specific user interface, which additionally records activity for a given mobility study (e.g. mode of transport, aim of transport, number of passengers etc.). Data is directly transmitted to a central server, allowing realtime monitoring of the audience to see if data is generated according to data campaign guidelines (cfr. Fig. 3.4). Messages can be sent to test persons and the diary interface can be changed on-the-fly to refine or even change the study.

It has been noted that in many travel behaviour studies, respondents do not always point out the right use of mode. In order to further minimize manual input from the user, automated classification of transport mode (foot, bike, car, public transport) is very useful. On a low level, this is typically performed based on a combination of accelerometer, GPS and/or magnetometer. Table 3.1 summarizes a number of studies on transportation mode classification based on smartphone data that have been described in literature. The table summarizes the main classification results (i.e. accuracy noted under performance) and the size of the dataset used. These results show that in general satisfactory classification performance is possible, although only limited tests have been reported. The majority of methods use a combination of ACM and GPS data, and the most common preprocessing is a combination of FFT transform and simple signal properties such as signal variance. Measured ACM and GPS data are collected in batches ranging from seconds to several minutes. These



**Fig. 3.4** (a) Distribution of individuals according to the number of unique days they used the wifi network (y-axis) vs. the timespan between the first and last time they used the network (x-axis). (b) Distribution of individuals according to their density, i.e. total time using the network divided by the time span between the first and last time using the network (y-axis) vs. the number of unique days they used the network (x-axis). The three visitor scenarios (S1, S2, S3) are highlighted

**Table 3.1** Comparison of techniques used for automated transport mode classification

Methods used	Performance	Sensors	Preprocessing	#P	#M	References
KNN, NB, DT, SVM, CT-HMM, DHMM	98.8	ACM, GPS	FFT, variance, energy	6	8	Reddy et al. (2010)
DT	82.1	ACM, GPS	FFT, variance	1		Manzoni et al. (2010)
GDA, SVM	93.8	ACM	FFT, variance, energy, mean	3	4	Nham et al. (2008)
KMC, NB, DT, SVM, CT-HMM, DHMM	93.6	ACM, GPS	DFT+CFS	16	5	Mun et al. (2009)
DT, LR, MLP, NB, SVM	95.6	ACM	FFT	6	4	Yang et al. (2010)
DT	N/A	ACM, GPS	Variance, mean, histogram, stop rate, device usage rate			Jonsson (2010)
NB+boosting	82	ACM, GPS, magnetometer	First order statistics	1	2	Frendberg (2011)
MLP	91.2	GPS	First order statistics		3	Gonzalez et al. (2008)

*GDA* Gaussian discriminant analysis, *KNN* K-nearest neighbours, *NB* Naive Bayes, *DT* Decision trees, *SVM* Support vector machine, *HMM* Hidden Markov model, *KMC* K-means clustering

#P stands for the number of people involved in the trial, #M represents the number of modi to be classified

batches are resampled to a uniform sampling rate and preprocessed to a frequency spectrum using fast fourier (FFT) and deriving associated frequency measures like variance and energy. Supervised classifier methods are trained using manual ground truth data.

### ***3.3.3 Characterising Urban Space Based on Population-Level Movement: Tourist Spotting Using Wi-Fi***

While in the previous paragraph people data are collected *explicitly*, here we discuss how this can be achieved *implicitly*. In this paragraph we present original work and results on characterizing urban space by considering population-level movement.

Previous work has considered population movement as a whole, or has focused on a particular subpopulation (e.g. tourists) to derive such a characterization. In the introduction to this paragraph examples using mobile phone traces were cited. An important limitation of these approaches is the coarse spatial granularity that does not allow for characterization of precise street-level locations. To achieve higher granularity, researchers are increasingly turning to alternative datasets. Analysis of user-generated content is becoming increasingly popular, for example using geo-tagged photos to extract “place” and “event” information (Rattenbury et al. 2007). Other work has considered granular Wi-Fi data (Calabrese et al. 2010; Kim and Kotz 2005) but so far limited to campus scale. Often, mobility analysis attempts to cluster locations based on similarity to each other in terms of volume of visitors. For instance, researchers have demonstrated a bottom-up approach to grouping locations into clusters that exhibit similar temporal mobility patterns in terms of volume of visits (Kim and Kotz 2005), and subsequently labels these clusters according to a tacit understanding of both the locations as well as the mobility patterns there (Calabrese et al. 2010).

Our work focuses on identifying short-term visitors to a city, and seeks to characterize urban space by contrasting these short-term visitors with long-term city-dwellers. It is expected that these two subpopulations exhibit sharp differences in their mobility patterns, because short-term visitors are new to the city and have not had time to adapt their behaviour.

This case study uses log data from a free and open municipal Wi-Fi network for the year 2010. By nature, this technology is more granular than cell phone technology, and works well for both indoor and outdoor spaces with a typical range of dozens of meters. The data originate from a network deployed in the city of Oulu, Finland, and consists of almost 1300 access points deployed in various parts of the city, but more densely in the downtown area and the university campus area. In particular, one data record exists for every time a device like a mobile phone or laptop associated or disassociated with the wireless network. In total the data consists of 7.8 million records for more than 82,000 devices, with all devices using the Wi-Fi network for a total of 200 million minutes, which is equivalent to 382

years. Finally, the population of the city (1500 km<sup>2</sup>) is 140,000, while the broader region (37,000 km<sup>2</sup>) has 385,000 residents.

To analyse these Wi-Fi traces one can rely on the identification of a particular segment of the population, in this case short-term visitors to the city. Previous work has shown that the use of temporal thresholds can identify visitors. For example, Girardin et al. (2008) define tourists as people whose activity is limited to a 30-day period in a city. The use of arbitrary temporal thresholds is the only way to passively identify “visitors”, simply because the term “visitor” is vague. For instance, the three following scenarios exemplify diverse instances of visitors:

- Scenario S1: a person visits a city for 5 days each month of the year
- Scenario S2: a person visits a city for 2 consecutive months during the year
- Scenario S3: a person visits a city for 5 consecutive days during the year

In this study we are interested in Scenario S3, i.e. understanding how short-term visitors behave differently from others. Hence, one can use a similar but more conservative approach as Girardin et al. (2008) and choose to focus on visitors who used the network for less than 10 days during 2010.

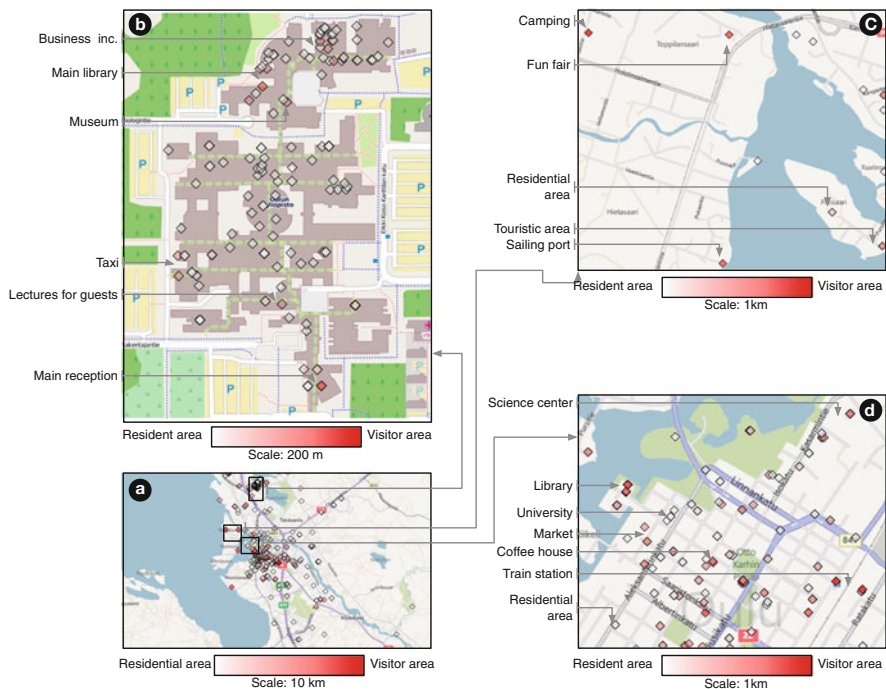
However, an initial analysis of the data suggests that this approach to defining visitors alone can be problematic due to the intermittent nature of Wi-Fi connectivity (Kim and Kotz 2005). It is quite possible that residents who live outside the Wi-Fi coverage area only use the network sporadically, especially if their daily routine does not involve a trip downtown. In Fig. 3.4a this intuition is validated by constructing a scatterplot where each dot represents a device from the dataset. One can observe that devices that used the network for a particular number of days during 2010 greatly varied in the timespan between the first and last time they used the network. For this reason a metric of “density” is developed, defined as the total time that a device used the network divided by the timespan of the device in the dataset (the time between its first and last sighting in the data). This way Fig. 3.4b can be constructed by plotting density versus the total number of days that a device used the network. This graph is annotated by indicating where the three scenarios described above fit, indicating that short-term visitors (Scenario S3) are in the top-left quadrant in Fig. 3.4b.

The next analysis stage characterizes various locations across the city by considering the differences in urban movement between visitors and residents. An initial approach was to characterize each location based on its popularity with visitors as proposed by Girardin et al. (2008). Careful inspection of these results, however, revealed that many places that were relatively popular with visitors are also popular amongst residents. A good example is the airport: the airport, along with other busy places, is popular with both visitors and residents. This particular issue has typically been ignored by researchers, either because the analysis did not focus on multiple segments of the population (Calabrese et al. 2011b; Calabrese et al. 2010; Kostakos et al. 2010; Quercia et al. 2011) or because there was no

comparison between segments of the population (Girardin et al. 2008; Kim and Kotz 2005; Rattenbury et al. 2007).

To address this issue a profile metric is developed to characterize the nature of a place by considering the relative popularity of a location with visitors and residents separately. To derive this metric one first calculates the “visitor\_power” and “resident\_power” of a location by normalizing the popularity of that location in relation to the most popular location for each group respectively. Then one may calculate the profile metric as  $visitor\_power / (visitor\_power + resident\_power)$ , ranging between 0 (“residential” area) and 1 (“visitor” area). This metric can then be used to characterize urban locations as shown in Fig. 3.5, where the airport is now profiled as a “visitor” area. An important advantage of this approach to location characterization is that it relies on an explicitly defined notion of human behaviour (in this case short-term visitors), while previous work has aimed to cluster locations that are similar to each other while being agnostic of human behaviour (Calabrese et al. 2010; Kim and Kotz 2005).

Figure 3.5 (left) shows the results of our characterization of a downtown area. The metric accurately profiles as resident areas a downtown university building



**Fig. 3.5** Using the profile metric one can inspect at a very granular level the whole city, but also specific locations: (left) parts of downtown, (right) a summer destination location



where mostly students are present. On the other hand, the market and the main library are profiled as visitor areas, as well as the science centre and the train station. Figure 3.5 (right) shows that the camping, beach and fun fair areas (which mostly operate in the summer) are profiled as visitor areas. We expected this to be the case, as these are locations that are likely to attract tourists during the summer months. It is also noteworthy that the touristic area with restaurants is indeed profiled as such, despite being within 250 m of a residential neighbourhood that is also correctly profiled.

On the campus area of the University all lecture halls and research lab areas are characterized as resident, obviously due to local students and staff occupying the areas. A handful locations are identified as visitor areas, namely the main reception of the university, a taxi rank, and the hall where lectures with outside guests are held. In addition, the main library with a substantial historical archive, the museum, and the business incubator are profiled as visitor areas. The profile metric is a surprisingly accurate match to the nature of these locations concerning visitors.

The work we presented in this case study shows a way to characterize human behaviour, and a way to characterize locations, both using implicitly sensed data. Both people and location characterization can be done automatically by the infrastructure. This means that locations can e.g. automatically infer their own profile, and network providers can tailor services for users based on their profile and the profile of their location. Patterns of urban mobility have previously been shown to be crucial in our understanding of urban space, and here we show techniques for extracting these patterns in a highly localized fashion.

More importantly, the method described here allows for analyses beyond the context of short-term visitors. This case study has explicitly referred to a “profile” metric because it can be used with any segmentation of the population. Instead of considering short-term visitors, the analysis can focus on males vs. females, teenagers vs. elders, locals vs. foreigners, or any other segmentation that can be reliably captured in the data. This approach to characterizing locations by considering human behaviour can provide a new way to look at cities, communities, and behaviour over time.

### **3.4 Observing Encounters and Activities Using Tagging**

Observing human activities using ubiquitous devices can be implemented using special sorts of tags. Bluetooth or Radio-frequency identification (RFID) tags, for example, provide cost-effective solutions for implementing sophisticated observation and tracking solutions. As outlined above, Bluetooth implemented in mobile phones provides for encounters between different devices by active scanning of the environment. The SDC Framework (Atzmueller et al. 2013) also allows to

collect various physical (e.g., Bluetooth) and also virtual sensor data (e.g., for accessing online social services) on mobile Android devices, for comprehensive ubiquitous data collection, location tracking and tagging. The Find-And-Connect (Xu et al. 2011) system also utilizes Bluetooth and passive RFID for obtaining locations of participants, and infers encounters based on the co-location of participants as a proxy of contacts between participants. However, no direct face-to-face contacts are measured. Essentially, such solutions—as well as passive RFID tags—enable the tracking of participants. However, given the range of interaction of Bluetooth devices, the detected proximity does not necessarily correspond to face-to-face contacts, see (Cattuto et al. 2010). As described below, the Sociopatterns collaboration, e.g., (Barrat et al. 2010) developed an active RFID-based proximity tag for detecting close-range and face-to-face proximity (1–1.5 m) as described below. An experiment correlating RFID-based proximity data with Bluetooth-encounter data collected using the SDC Framework is described in (Atzmueller et al. 2013). Another approach for observing human activities is the Sociometric Badge (<http://hd.media.mit.edu/badges>). It records more details of human activities and interactions than only using the SocioPatterns RFID tags, but requires significantly larger devices compared to these tags.

### ***3.4.1 RFID Proximity Tagging***


The SocioPatterns collaboration developed an infrastructure that detects close-range and face-to-face proximity (1–1.5 m) of individuals wearing special RFID (proximity) tags with a temporal resolution of 20 s (Cattuto et al. 2010). The technical innovation of these tags is their ability to detect the proximity of other tags within a range of up to 1.5 m. Since the human body blocks RFID signals, face-to-face contacts can then be detected by tag proximity. Essentially, each proximity tag sends out two types of RFID-signals, proximity signals and tracking signals. A proximity signal is used for contact sensing, which is achieved using signals with very low radio power levels. The proximity tag further sends out tracking signals in four different signals strengths to specialized RFID readers. These tracking signals are used to transmit proximity information to a central server and for determining the position of each conference participant. Depending on the signal strength, the range of a tracking signal inside a building is up to 25 m. RFID readers measure a received signal's strength, the ID of the reporting tag and the IDs of all RFID tags in proximity. For more information about the SocioPatterns proximity tags see (Barrat et al. 2010), and the web-sites of SocioPatterns (<http://www.sociopatterns.org>) and OpenBeacon (<http://www.openbeacon.org/>).

### 3.4.2 *Observing Encounters and Activities*

Experience showed (Atzmueller et al. 2014) that people are more motivated to wear such tags—in particular over a longer period of days or even weeks—if they gain a personal benefit. The SocioPatterns platform has been utilized in different scenarios, for example, for tracking conference attendees (Alani et al. 2009). It was also deployed in other applications, such as healthcare environments (Isella et al. 2011), schools (Stehle et al. 2011) and museums (Cattuto et al. 2010). Barrat et al. (2010) analyze social dynamics of conferences focusing on the social activity of conference participants in those experiments. They analyze, for example, their activity in social web platforms like Facebook, Twitter and other social media together with status and their research seniority.

In addition, the SocioPatterns RFID tags are also used for the Conferator (Atzmueller et al. 2011) and MyGroup (Atzmueller et al. 2012a) systems, which enable enhanced social networking for conference participants and working groups, respectively. Both systems are built using the Ubicon software platform, see (Atzmueller et al. 2014). It aims at enhancing ubiquitous and social networking, and enables the observation of physical and social activities. For more details, we refer to Chap. 6. Conferator and MyGroup have been applied at a number of events, for example, at the LWA 2010, LWA 2011 and LWA 2012 conferences of the German association of computer science, and at the ACM Hypertext 2011.

In this context, Atzmueller et al. (2012b) analyze the interactions and dynamics of the behavior of participants at conferences. Figure 3.6 shows an example of a temporal community analysis of the LWA 2010 conference. The figure shows overlapping communities for different minimal time thresholds of the contact network ( $t = 0, 1, 5, 10$ ) grounded by special interest groups (ABIS, IR, KDML, WM) present at the conference. The figure shows, that the communities tend to focus more on the special interest groups with an increasing minimum conversation length threshold, i.e., the communities start with a mixture of different interest groups, but concentrate more and more on special sub-communities with an increasing time threshold. The connection between research interests, roles and academic jobs of conference attendees is further analyzed in Macek et al. (2012). Furthermore, the predictability of links in face-to-face contact networks and additional factors also including online networks has been analyzed by Scholz et al. (2012, 2013).



**Fig. 3.6** Exemplary community detection results (overlapping communities) colored according to their special interest track distributions (Atzmueller et al. 2012). The figure shows the different special interest groups (ABIS, IR, KDML and WM). The x-axis depicts the sizes of the groups (and its subgroups, i.e., communities). From (a) to (d) the minimal conversation length threshold increases from 0 to 10 min, focusing on face-to-face contacts having at least that length

### 3.5 Conclusion

We have given an overview of technological solutions that allow to observe human activity. Observation by scanning, location-enabled devices and tagging have been discussed. Where smart camera networks potentially offer the most detailed analysis of human activity, the scaling of these systems to large environments currently remains limited due to the limits on the performance of the current algorithms for recognition of individuals for large databases. When observing humans outside vehicles, in multimodal transport or in indoor situations, scanning

based on Bluetooth, WiFi, RFID tags or smartphones all offer a possible solution. It depends on the level of detail required, the size of the target crowd and the effort possible for activating the crowd for data generation that will determine the optimal solution for a successful data campaign.

## References

- Alani, H., Szomszor, M., Cattuto, C., den Broeck, W.V., Correndo, G., Barrat, A.: Live social semantics. In: *Proceedings of the International Semantic Web Conference*, pp. 698–714. (2009)
- Asakura, Y., Hato, E.: Tracking survey for individual travel behaviour using mobile communication instruments. *Transp. Res. C* **12**, 273–291 (2004)
- Asmundsdottir, R., Chen, Y., van Zuylen, H.J.: Dynamic origin-destination matrix estimation using probe vehicle data as a priori information. In: Barcelo, J., Kuwahara, M. (eds.) *Traffic Data Collection and Its Standardization*, pp. 89–108. Springer, New York (2010)
- Atzmueller, M., Benz, D., Doerfel, S., Hotho, A., Jäschke, R., Macek, B.E., Mitzlaff, F., Scholz, C., Stumme, G.: Enhancing social interactions at conferences. *Inf. Technol.* **53**(3), 101–107 (2011)
- Atzmueller, M., Becker, M., Doerfel, S., Kibanov, M., Hotho, A., Macek, B.-E., Mitzlaff, F., Mueller, J., Scholz, C., Stumme, G.: Ubicon: observing social and physical activities. In: *Proceedings of the 4th IEEE International Conference on Cyber, Physical and Social Computing (CPSCom)*, 2012
- Atzmueller, M., Doerfel, S., Mitzlaff, F., Hotho, A., Stumme, G.: Face-to-face contacts at a conference: dynamics of communities and roles. In: *Modeling and Mining Ubiquitous Social Media*, vol. 7472. Springer, Heidelberg, 2012
- Atzmueller, M., Hilgenberg, K.: Towards capturing social interactions with SDCF: an extensible framework for mobile sensing and ubiquitous data collection. In *Proceedings of the 4th International Workshop on Modeling Social Media*, ACM, 6 (2013)
- Atzmueller, M., Becker, M., Kibanov, M., Scholz, C., Doerfel, S., Hotho, A., Macek, B.-E., Mitzlaff, F., Mueller, J., Stumme, G.: Ubicon and its applications for ubiquitous social computing. *New Rev. Hypermedia Multimedia* **20**(1), 53–77 (2014)
- Atzori, L., Iera, A., Morabito, G.: The internet of things: a survey. *Comput. Netw.* **54**(15), 2787–2805 (2010)
- Barrat, A., Cattuto, C., Szomszor, M., den Broeck, W.V., Alani, H.: Social dynamics in conferences: Analyses of data from the live social semantics application. In: *Proceedings of the International Semantic Web Conference*, LCNS, vol. 6497, pp. 17–33. (2010)
- Bredereck, M., Jiang, X., Korner, M., Denzler, J.: Data association for multi-object tracking-by-detection in multi-camera networks. In: *Proceedings of the Sixth International Conference on Distributed Smart Cameras (ICDSC2012)*, 2012
- Buch, N., Velastin, S.A., Orwell, J.: A review of computer vision techniques for the analysis of urban traffic. *IEEE Trans. Intell. Transp. Syst.* **12**(3), 920–939 (2011)
- Bullock, D., Haseman, R., Wasson, J., Spittler, R.: Anonymous bluetooth probes for measuring airport security screening passage time: the indianapolis pilot de-ployment. In: *Transportation Research Board 89th Annual Meeting*, 2010
- Calabrese, F., Di Lorenzo, G., Liu, L., Ratti, C.: Estimating origin-destination flows using mobile phone location data. *IEEE Pervasive Comput.* **10**, 36–44 (2011a)
- Calabrese, F., Ratti, C., Colonna, M., Lovisolo, P., Parata, D.: Real-time urban monitoring using cell phones: a case study in Rome. *IEEE Trans. Intell. Transp. Syst.* **12**(1), 141–151 (2011b)
- Cattuto, C., den Broeck, W.V., Barrat, A., Colizza, V., Pinton, J.F., Vespignani, A.: Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLoS One* **5**(7) (2010)

- Eagle, N., Pentland, A.: Reality mining: sensing complex social systems. *Personal Ubiquit. Comput.* **10**(4), 255–268 (2006)
- Frendberg, M.: Determining transportation mode through cell phone sensor fusion. PhD thesis, Massachusetts Institute of Technology, Cambridge (2011)
- Girardin, F., Calabrese, F., Fiore, F.D., et al.: Digital footprinting: uncovering tourists with user-generated content. *IEEE Pervasive Comput.* **7**, 36–43 (2008)
- Gong, H., Chen, C., Bialostozky, E., Lawson, C.T.: A GPS/GIS method for travel mode detection in New York city. *Comput. Environ. Urban Syst.* **36**(2), 131–139 (2011)
- Gonzalez, P., Weinstein, J., Barbeau, S., Labrador, M., Winters, P., Georggi, N.L., Perez, R.: Automating mode detection using neural networks and assisted gps data collected using gps-enabled mobile phones. In: 15th World Congress on Intelligent Transportation Systems, 2008
- Hato, E.: Development of behavioral context addressable loggers in the shell for travel-activity analysis. *Transp. Res. C* **18**, 55–67 (2010)
- Isella, L., Romano, M., Barrat, A., Cattuto, C., Colizza, V., Van den Broeck, W., Tozzi, A.E.: Close encounters in a pediatric ward: measuring face-to-face proximity and mixing patterns with wearable sensors. *PLoS One* **6**(2), e17144 (2011)
- Jonsson, F.: Determining transportation mode in mobile phones using human agent movements. *USCCS* 73 (2010)
- Kim, M., Kotz, D.: Modeling users' mobility among WiFi access points, Workshop on Wireless traffic measurements and modeling, pp. 19–24. *USENIX Association* (2005)
- Kostakos, V., O'Neill, E.: Cityware: Urban computing to bridge online and real-world social networks. In: Foth, M. (ed.) *Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City*, pp. 195–204. IGI Global, Hershey/London (2008)
- Kostakos, V., O'Neill, E., Penn, A., et al.: Brief encounters: Sensing, modeling and visualizing urban mobility and copresence networks. *ACM TOCHI* **17**, 2 (2010)
- Kruegle H.: *CCTV Surveillance: Video Practices and Technology*. Elsevier Butterworth-Heinemann, Burlington/Oxford (2011)
- Macek, B.E., Scholz, C., Atzmueller, M., Stumme, G.: Anatomy of a conference. In: *Proceedings of the 23rd ACM conference on Hypertext and Social Media*, pp. 245–254. *ACM* (2012)
- Malinovskiy, Y., Wu, Y.J., Wang, Y., Lee, U.K.: Field experiments on bluetooth-based travel time data collection. In: *Transportation Research Board 89th Annual Meeting*, 2010
- Manzoni, V., Maniloff, D., KloECKl, K., Ratti, C.: Transportation mode identification and real-time CO2 emission estimation using smartphones. Technical report, Massachusetts Institute of Technology, Cambridge (2010)
- Morbee, M., Tessens, L., Aghajan, H., Philips, W.: Dempster-Shafer based multi-view occupancy maps. *Electron. Lett.* **46**(5), 341–343 (2010)
- Mun, M., Reddy, S., Shilton, K., Yau, N., Burke, J., Estrin, D., Hansen, M., Howard, E., West R., Boda, P.: PEIR, the personal environmental impact report, as a platform for participatory sensing systems research. In: *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services*, pp. 55–68. *ACM, New York* (2009)
- Nham, B., Siangliulue, K., Yeung, S.: Predicting Mode of Transport from iPhone Accelerometer Data. Technical report, Stanford University (2008)
- O'Neill, E., Kostakos, V., Kindberg, T., Penn, A., Fraser, D. S., Jones, T.: Instrumenting the city: Developing methods for observing and understanding the digital cityscape. In: *UbiComp 2006: Ubiquitous Computing*, pp. 315–332. Springer, Berlin Heidelberg (2006)
- Pentland, A.: Looking at people: Sensing for ubiquitous and wearable computing. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(1), 107–119 (2000)
- Quercia, D., Di Lorenzo, G., Calabrese, F., Ratti, C.: Mobile phones and outdoor advertising: measurable advertising. *IEEE Pervasive Comput.* **10**, 28–36 (2011)
- Rattenbury, T., Good, N., Naaman, M.: Towards automatic extraction of event and place semantics from flickr tags. In: *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Amsterdam, 23–27 July 2007
- Reddy, S., Mun, M., Burke, J., Estrin, D., Hansen, M., Srivastava, M.: Using mobile phones to determine transportation modes. *ACM Trans. Sens. Netw. (TOSN)* **6**(2), 13 (2010)

- Reinthal, M., Zajicek, J.: Real time route analysis based on floating car technology. In: 18th IASTED International Conference Modelling and Simulation, pp. 609–612. Montreal, Canada (2007)
- Scholz, C., Atzmueller, M., Stumme, G.: On the predictability of human contacts: influence factors and the strength of stronger ties. In: Proceedings of the Fourth ASE/IEEE International Conference on Social Computing (SocialCom), IEEE Computer Society, Boston, MA, 2012
- Scholz, C., Atzmueller, M., Barrat, A., Cattuto, C., Stumme, G.: New insights and methods for predicting face-to-face contacts. In: Proceedings of the 7th International AAAI Conference on Weblogs and Social Media (ICWSM'13), 2013
- Stehle, J., Voirin, N., Barrat, A., Cattuto, C., Isella, L., Pinton, J.F., Quaghiotto, M., den Broeck, W.V., Regis, C., Lina, B., Vanhems, P.: High-resolution measurements of face-to-face contact patterns in a primary school. *CoRR* (2011). abs/1109.1015
- Taale, H., Hoog, A.D., Smulders, S., Tool, O.: The results of a Dutch experiment with floating car data. In: Control in Transportation Systems 2000: A Proceedings of the Volume from the 9th IFAC Symposium, vol. 1. Braunschweig, Germany, 2001
- Torp, K., Lahrman, H.S.: Floating car data for traffic monitoring. In: ITS at the Crossroads of European Transport, ERTICO—ITS Europe, 2005
- Tsui, A., Shalaby, A.: Enhanced system for link and mode identification for personal travel surveys based on global positioning systems. *Transp. Res. Rec. J. Transp. Res. Board* **1972**(1), 38–45 (2006)
- Versichele, M., Neutens, T., Delafontaine, M., Van de Weghe, N.: The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities. *Appl. Geogr.* **32**(2), 208–220 (2012)
- Vlassenroot, S., Gillis, D., Bellens, R., Gautama, S.: The use of smartphone applications in the collection of travel behavior data. In: 9th ITS European Congress, Proceedings, ERTICO, 2013
- Weinzerl, J., Hagemann, W.: Automatische Erfassung von Umsteigern per Bluetooth Technologie. *Nahverkehrspraxis* **3**, 18–19 (2007)
- Xie, X., Grünwedel, S., Jelaca, V., Niño Castañeda, J., Van Haerenborgh, D., Van Cauwelaert, D., Van Hese, P., et al.: Learning about objects in the meeting rooms from people trajectories. In: 2012 Sixth International Conference on Distributed Smart Cameras (ICDSC), 2012
- Xu, B., Chin, A., Wang, H., Chang, L., Zhang, K., Yin, F., Wang, H., Zhang, L.: Physical proximity and online user behavior in an indoor mobile social networking application. In: Proceedings of the 4th IEEE International Conference on Cyber, Physical and Social Computing (CPSCom2011), 2011
- Yang, Y., Toida, T., Hong, C.: Transportation prediction using build-in triaxial accelerometer in cell phone. In: International Conference on Business Information, Bai, 2010