Collaborative Sensing for Urban Transportation

Sergio Ilarri[1], Ouri Wolfson[2], and Thierry Delot[3]

[1]University of Zaragoza, Spain, silarri@unizar.es
[2]University of Illinois, Chicago, USA, wolfson@cs.uic.edu
[3]University of Valenciennes, France, Thierry.Delot@univ-valenciennes.fr

Abstract

In this paper, we overview the current status of research in the field of collaborative sensing for urban transportation. We first classify the different types of sensors that can be relevant in collaborative urban sensing and then we analyze three different topics: parking spaces, traffic, and trajectories. We discuss issues regarding the sensing and data sharing approaches as well as relevant use cases. Finally, we identify some research challenges and trends.

1 Introduction

A key element for urban informatics [1] is the widespread availability of sensors of different types, which can capture a wide range of environmental conditions and information of the surroundings of the user. Besides traditional wireless sensor networks, where sensors are statically deployed over a fixed area to measure certain values, mobile and urban computing scenarios open up new opportunities for more flexible and dynamic approaches. Collaborative sensing (also called mobile crowdsensing or participatory sensing) [2, 3] implies using sensors in a cooperative way in order to obtain an overall and more complete perspective of the environment. Specifically, the new scenarios support the use of *mobile sensors*. On the one hand, users with their mobile devices have become an important source of sensor data [4]. On the other hand, vehicles can also carry a variety of sensors ("Today's luxury cars have more than 100 sensors per vehicle" [5]).

Collaborative sensing can provide key benefits in urban transportation, contributing to higher travel efficiency, safety, and reduced pollution, through innovative applications that benefit from the data sensed collaboratively. These benefits apply to private transportation as well as public transportation. The collaborative aspect remarks the idea that the role of people is particularly important. For example, the work presented by Campbell et al. adopts a human-centric (or people-centric) view of urban sensing [6] , which emphasizes the significance of the attributes of people, their surroundings, and their interactions with the environment.

In this paper, we survey the state of the art in collaborative sensing for urban transportation. In Section 2 we overview the different types of sensors that could be of relevance. In Section 3 we focus on the problem of sensing and sharing information about parking spaces, as a prime example of a scarce resource in urban transportation (bike stations could be another example). In Section 4 we analyze the problem of traffic estimation

Copyright 0000 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.
Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

from a cooperative sensing perspective. In Section 5 we tackle the management of trajectories. Finally, in Section 6 we summarize some existing challenges. Figure 1 shows an overview of the main topics covered.

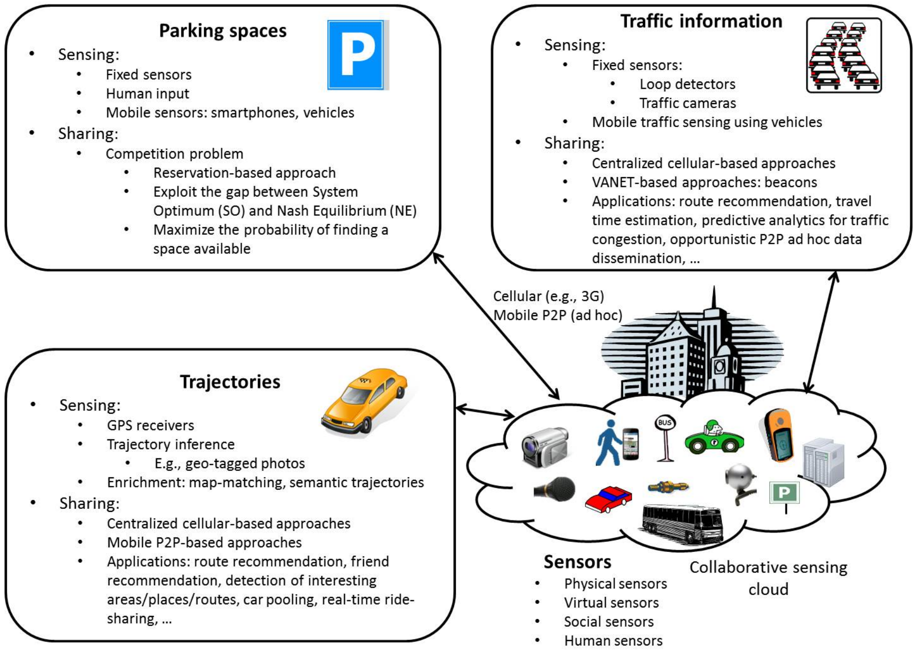

Figure 1: Main topics in collaborative sensing for transportation

2 Types of Sensors

According to the abstraction level of the information they manage, several types of sensors can be distinguished:

- *Physical sensors*. These are traditional *hardware sensors*, which are able to directly capture raw data from the environment. Sensing capabilities can be embedded in existing mobile devices (e.g., smartphones) or as stand-alone measuring devices that communicate their data to other devices or computers. As an example, mobile devices are being equipped with an increasing number of built-in sensors: GPS, microphone, camera, ambient light sensor, accelerometer, gyroscope, compass, proximity sensor, temperature, etc. These mobile devices can be carried by people or be integrated in vehicles (cars, taxis, buses, etc.).

- *Virtual sensors*. These are *software sensors*, which do not directly correspond to physical sensors. Instead, they usually provide higher-level data that are obtained by combining the output of several sensors. A virtual sensor [7] is an abstraction that aggregates data provided by different (and possibly heterogeneous) sensors to compute a virtual measurement (e.g., a virtual position sensor that transparently uses different positioning mechanisms, as needed, to locate the user).

- *Social sensors*. These are sensors that provide data based on information extracted from the social media: social networks (e.g., Facebook, Foursquare, Flickr), blogs, microblogs (e.g., Twitter), etc. For example, Albakour et al. [8] exploit microblogs to detect events in the vicinity, and Sakaki et al. [9] extract data from tweets to infer heavy traffic conditions for drivers. Another prominent effort is the EU FP7 project *SocialSensor* [10].

- *Human sensors*. In the extreme, we can consider humans as sensors, as they themselves can provide interesting measures through explicit cooperation. Humans can use their own senses and provide information explicitly, or manage the sensors that they have available in specific ways to collaborate in the measuring task. As an example, humans can provide *volunteered geographic information (VGI)* [11]. As another example, the idea of *spatial crowdsourcing* [12, 13, 14] is a hot topic. Drivers can also provide information about parking spaces (see Section 3). Although in some cases social sensors could also be considered as human sensors (e.g., humans can report what they see by sending tweets), we make a distinction by requiring explicit (not indirect or casual) cooperation from people in the case of human sensors.

The previous classification is to some extent similar to the one provided for location sensors in the work by Indulska et al. [15]. According to that work, physical location sensors provide the location of a physical device, virtual location sensors extract information from the virtual space (software applications, operating systems, and networks), and logical location sensors exploit virtual and physical location sensors to infer physical locations. On the other hand, Islam et al. [16] distinguish between *hard sensors* (physical sensors) and *soft sensors* (which actually represent trigger conditions, such as the reception of an email or a new entry in the user's calendar).

3 Parking Spaces

A key element to improve urban transportation is to provide a suitable management of parking spaces. Among the inconveniences of searching for parking, we can highlight that it is usually a time-consuming task, stressful and frustrating for drivers, and it obviously contributes to increasing traffic and pollution. A well-known report by Shoup [17] quantifies the yearly cost of parking in a commercial district next to the campus of the University of California in Los Angeles (Westwood Village): 47000 gallons of gasoline, 730 tons of CO_2 emissions, and 95000 hours of drivers' time. According to a study in France [18], searching for parking represents between 5% and 10% of the traffic in urban areas and it can rise up to 60% in the case of small streets. As a final example, a study performed in the district of Schwabing in Munich (cited in [19]) indicates that about 44% of the traffic is searching for an available parking space. These and many other reports emphasize the importance of promoting an efficient management of parking spaces for sustainable transportation in cities.

For some types of parking, such as parking garages and paid parking lots, an existing infrastructure can accurately monitor in real-time the occupancy of parking spaces and even provide effective reservation and payment methods. However, in the general case, such as in the case of free street parking, a fixed sensor-based infrastructure to detect the occupancy of parking spots (a popular example is SFPark in San Francisco, see `http://sfpark.org/`, which covers 6000 metered spaces and 12250 spaces in parking garages [20]) usually does not exist or it is considered expensive to be deployed at a large scale (e.g., see [21, 22, 23, 24]). Therefore, collaborative parking sensing is a promising avenue to explore.

It should be noted that a variety of methods could be used to capture information about the availability of a parking space. Besides traditional approaches using fixed sensors (e.g., SFPark), smartphone-based techniques are being developed to automatically detect parking and unparking events without user intervention [22, 23, 24, 25, 26]. Other approaches consider the vehicles themselves as mobile sensors that detect surrounding parking spaces [21]. Finally, the driver himself or herself could also provide the information manually (as illustrated in the work by Hoh et al. [27] or by a variety of ephemeral applications such as Google's Open Spot, which is no longer available).

Once the information about parking spaces is collected, the problem is how to effectively share it. Parking spaces represent a scarce resource, and therefore if information about their availability is disseminated freely among the potentially interested vehicles there is a high risk of generating a competition among them to try to get the same parking space. There is a trade-off between competition and information dissemination [28], such that eventually the time to find an available parking space could actually increase as compared to the situation of blind search where no information is provided. Different approaches have been proposed to manage this problem:

- Some approaches have proposed reservation protocols to guarantee the use of the parking space to only a specific vehicle (e.g., the *Centralized Assisted Parking Search* approach [28]). They rely on an infrastructure for reservation, payment and monitoring of the parking spaces, and so their use is limited to the case of controlled parking facilities. An exception is the reservation protocol for VANETs (Vehicular Ad Hoc Networks) proposed by Delot et al. [29], which uses the term *reservation* in a metaphorical way. What it actually does is to provide the information about an available parking space to a single driver); without the use of an infrastructure, ensuring that no other vehicle can take up the "reserved" space is not possible.

- Other approaches exploit the gap between the System Optimum (SO) and Nash Equilibrium (NE) in matching vehicles and parking slots [30]. SO is more efficient (i.e., it involves minimum total travel-time to slots), but an uncoordinated system of drivers acting selfishly settles into an NE state. One approach to bridge the gap is to price the slots such that an NE matching when considering travel-time and price is the same as an SO matching when considering travel-time alone [31]. Another approach involves payments among drivers looking to park, through a platform, to convert one matching to another [31, 32]. In other words, through payments, selfish drivers are incentivized to behave in a manner that is globally and environmentally responsible. Interestingly, payments raise the issue of truthful location disclosure by a driver, and Vickrey-Clarke-Groves (VCG) mechanisms from economics can be adapted to incentivize drivers to be truthful [32].

- Some approaches try to maximize the probability of finding an available parking space. For example, one alternative is to compute a route that goes through all the available parking spaces [33]. As another example, the Gravity-based Parking Algorithm (GPA) [30] is a gravitational model heuristic where each parking space applies an attraction force on a vehicle searching for a parking space. The magnitude of the force is inversely related to the distance between the space and the vehicle. The force-vectors on the vehicle are added, thus areas with a higher density of parking spaces become more attractive. In the approach proposed by Klappenecker et al. [34], parking lots periodically disseminate certain status parameters (such as their capacity and the number of occupied parking spaces) that vehicles can use to estimate the probability of finding there an available parking space at the time of arrival. Other proposals highlight the interest of guiding drivers towards areas where the probability to find an available parking space is high [35], instead of towards a specific parking space. Zekri et al. [36] present an approach that aggregates information about available parking spaces to extract general knowledge about their spatio-temporal availability.

Although most studies have focused on effectively exploiting availability data about parking spaces to minimize the parking time, the communication overhead should not be overlooked. As an example, a survey concerning the Montreal Region [37] (an area of 5500 km^2) reports 80500 parking events per day during the year 2003. As another example, from the analysis of a dataset of SFPark we conclude that an average of 56.8 records (events) per city block are generated each day in San Francisco; the raw data of each record has a size of 129.7 bytes on average (the names of the blocks have different lengths) which can be reduced (filtering out unnecessary or redundant information) for later processing to an average of 33 bytes of useful data per record. If a high number of messages about parking and unparking events are exchanged among vehicles with no control, this could easily lead to network overload. If we imagine a country or world-wide cloud-based system, the volumes of data in such a system may be particularly significant.

4 Traffic Information

For the collection of traffic information, different methods have been traditionally used. For example, for speed or traffic density estimation static devices and sensors such as single and double loop detectors or traffic surveillance cameras can be exploited. In the last years, rather than relying on sensing approaches based on the use of a fixed infrastructure, which is usually an expensive and non-flexible approach with limited coverage (e.g., see [38]), other alternatives propose the use of vehicles as mobile sensors.

Thus, *probe vehicles* (e.g., taxis and buses) could collect traffic data in a city. For example, they could periodically report their speed and location to a central server (using cellular communication) for later aggregation [38]. This collaborative traffic sensing is a promising approach, but it also faces important challenges: the data are sensed opportunistically as the vehicles move through the city, and therefore the data collected may exhibit significant gaps in time and space, as the spatio-temporal distribution of the vehicles is not uniform (e.g., compressive sensing could be used to deal with missing values [38]). Vehicles can use different types of sensors to obtain the relevant information; for example, VTrack [39] exploits WiFi location samples to estimate the location of the user (with less energy consumption than in a GPS-based solution) and identify travel-time patterns along the road segments. Besides probe vehicles, *floating cars* (such as patrol cars for surveillance) can also play the role of mobile sensors [42] for traffic estimation: the path of these floating vehicles could be controlled and adjusted according to the traffic monitoring requirements.

The wireless bandwidth used for the transmission of traffic information is a resource to minimize [40, 41]. For example, Ayala et al. [41] proposed a *flow-based update policy* where cars use a transmission probability computed based on the number of messages that the server expects to receive from the vehicles in each road segment in order to ensure a desired accuracy for the average speeds computed. They performed experiments based on data generated by a sensor at the end of a road segment on highway 190 in Chicago from 6AM to 8AM on May 22nd, 2007 (Tuesday). In the experiments presented, fewer that 500 messages are required by the proposed policy, whereas the number of messages goes up to about 3000 with an alternative *information cost based policy* (based on the existing trade-off between the communication cost and the data uncertainty) and about 5500 for a *deterministic policy* (a vehicle transmits its velocity to a server if the difference between the broadcasted velocity received from the server and the measured velocity exceeds a certain threshold), considering a data collection period of 300 seconds and a velocity threshold of 1 mph for the information cost based policy and deterministic policy; the benefits of the proposed policy are also shown for other combinations of thresholds.

Whereas traditional mobile traffic sensing approaches using vehicles are centralized and usually use cellular communications, other novel solutions have been proposed in the field of VANETs, which exploit ad hoc communications in a distributed way. In the proposal presented by Sanguesa et al. [43], vehicles in urban areas use information of the beacons received from other neighboring vehicles, as well as features extracted from digital maps, to estimate the density of vehicles in their neighborhood; in the experimental evaluation performed, the authors infer that the optimal time period to estimate the density in vehicular environments is 30 seconds. To offer an overall picture rather than just a local density estimation, some proposals advocate exploiting a complementary fixed-support infrastructure. Thus, the V2X-D architecture [44] combines vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications to estimate the traffic density, based on the number of beacons received by vehicles and fixed support nodes on the roads (Road Support Stations or RSUs) and the features of the road maps. Another recent VANET-based approach for speed estimation is proposed by Yang et al. [45].

It is also interesting to mention that it is possible to combine information provided by (fixed or mobile) traffic sensors with other data sources. For example, several approaches [46, 47, 48] combine data provided by traditional monitoring equipment with social media feeds to obtain a better perception and try to infer the causes of the traffic situation.

The estimation of traffic density and travel times has a range of potential benefits. Providing drivers with estimated travel times would help them to make more informed decision and take the best routes [49], even in multimodal settings involving both public and private transportation. Moreover, traffic density has a significant impact on the performance of dissemination protocols for VANETs, and therefore the knowledge of the traffic density could be exploited to have a more robust data dissemination or routing protocol (e.g., see [50]). Traffic awareness also implies the detection of hazards, such as traffic jams, accidents, or a malfunctioning bus line.

Besides estimating the current situation, once traffic data have been collected, it is also interesting to analyze them to learn something useful that can help to improve traffic-based applications, such as travel recommendation systems, even in the absence of real-time information. Predictive analytics, which implies extracting information from data in order to predict future outcomes and trends, is a hot topic currently. Its potential application to the field of transportation, particularly exploiting traffic information, is very promising. For example, Chong et al. [51] present a use case that predicts traffic congestion by using a collaborative analytics system proposed by the authors. Xu et al. [52] describe a time-series approach to prediction for processing moving objects queries, and Min et al. [53] enhance it by exploiting spatial and temporal interactions and considering the road network characteristics (the possible congestion or not of a link influences the traffic flow of others).

As an example application, Waze (`http://www.waze.com`) is a popular system allowing a community of drivers to share information (e.g., traffic information). It captures contextual information (e.g., location, speed) and aggregates the data collected to deliver real-time information about the surroundings of the user (e.g., location and size of traffic jams). Another example is the exchange of multimedia data about traffic [54]; if a 5 second video or voice recording is captured by each vehicle every minute and disseminated, the data volumes generated and exchanged may be significant.

5 Trajectories

Trajectory-data management is another relevant topic in the field of urban transportation. We can identify three basic tasks regarding the collection, processing, and use of trajectories:

- *Capturing trajectory information*. The use of GPS is now widespread and many mobile devices (such as modern smartphones) are equipped with a GPS receiver. With such a positioning mechanism (or using other similar satellite-based techniques, such as the upcoming European *Galileo*), the trajectory of the device can be easily captured as a sequence of location data points. Moreover, less-precise positioning techniques (e.g., network-based positioning [55]) can be applied outdoors if a GPS receiver is not available or if the device is in a covered area with fewer than four satellites in view. Finally, in some cases it could be possible to infer trajectories, such as popular travel paths, from other sources like geo-tagged pictures [56].

- *Enriching trajectories*. Raw location data is sometimes of little use. Therefore, the GPS locations have to be cleaned and corrected. A common preprocessing step, in the case of vehicles, is applying *map-matching* [57], which is based on the assumption that vehicles move along roads and that their movement is smooth. This is not only a cleaning technique that allows the correction of GPS errors, but it also leads to data enrichment, as it allows labeling raw locations with the corresponding streets and roads. This enrichment can be further enhanced, going from *raw trajectories* to *semantic trajectories* [58], which not only encode data points but also higher-level semantic information. For example, semantic annotations could represent the type of movement (walking, running, driving, etc.), the transportation means used (private car, bus, taxi, etc.) [59], or any other feature of interest, such as the type of business visited [60]. For example, trajectory data enriched with the transportation means could be useful for other users to discover new multimodal routes using both public and private transportation. According to Ilarri et al. [61], different mobility dimensions and levels of granularity could be considered, depending on the needs.

- *Sharing trajectory data*. A number of websites have been developed to support users in the task of sharing and searching trajectory data (e.g., ShareMyRoutes and Bikely). These sites allow users to share and find information about interesting touring routes, bike routes, hiking routes, etc. Similarly, there are also many similar applications for smartphones, such as RouteShoot, which captures both a GPS trace and a video of the surroundings of the trajectory. These applications and web sites are based on a centralized architecture, where a server receives the information through the Internet (the mobile devices use cellular communications such as 3G) and can serve trajectory requests from other interested users. Nevertheless, we could also envision advanced applications where the trajectory data are also shared using peer-to-peer (P2P) ad hoc wireless communications in an opportunistic way. As an example, although it is not a pure P2P architecture because it exploits an infrastructure of access points, the *Shared-Trajectory-based Data Forwarding Scheme* (STDFS) [62] for VANETs is based on the estimation of an encounter graph that is exploited for packet forwarding.

The previous tasks would make trajectory data available in an appropriate form. Then, it is also important to develop techniques that can efficiently and effectively exploit the trajectory information. Trajectory data can be collected and mined to learn and extract interesting information. For example, the T-Drive system [63] mines historical GPS trajectories of taxis (taxi drivers are assumed to be experts in finding the fastest routes) in order to recommend appropriate time-dependent driving directions. Ying et al. [64] advocate the use of a semantic similarity measure between trajectories as the basis of a friend recommender system. Analyzing trajectories to detect interesting areas and popular routes is another key application (e.g., see [65]). The information about popular areas and routes can be exploited in recommendation systems, traffic and service planning, and route finding. Trajectories can also be exploited for car pooling, by matching users with similar profiles based on historical information [66], or even for real-time ride-sharing [67]. Finally, as mentioned above, trajectories provide very useful information for opportunistic data forwarding in delay-tolerant networks (e.g., see [62]). Depending on the specific scenario considered and the purpose of the trajectory-data management, huge volumes of trajectory data may need to be processed; for example, in the study by Giannotti et al. [68] the authors consider a data set of around 17000 cars performing around 200000 travels over a week in the city of Milan (Italy), as well as a dataset of around 40000 cars performing around 1500000 travels over 5 weeks in the region surrounding the city of Pisa (Italy).

6 Some Challenges and Future Trends

Although the wide availability of sensors offers very interesting opportunities for urban sensing in transportation, there are also challenges that need to be tackled, such as:

- *Encouraging cooperation among the users*. A significant research topic is focused on providing incentives for users to cooperate. For example, Yang et al. [69] propose both a platform-centric model (the system itself provides a reward to the participating users) and a user-centric model (based on auctions, the users demand a specific reward in exchange of their sensing services). Another interesting related work is provided by Sugiyama et al. [70], who analyze incentive mechanisms in the context of data acquisition and distributed computing applications. Although these are significant research contributions, we believe that more effort is needed to propose specific incentive mechanisms for practical scenarios and evaluate their real impact on the behavior of the users. The use of *gamification* to engage users should also be further explored.

- *Management of the data quality*. With opportunistic sensing it is difficult to guarantee, or even estimate, a specific degree of quality for the data provided. This could be alleviated through the development of techniques that estimate the expected level of cooperation and the reliability of the different contributors. With this type of information available, it would be possible to quantify the accuracy, completeness, and freshness of the data. Guaranteeing a certain data quality could also require encouraging further cooperation among the users, although best-effort (rather than optimal) approaches should be expected.

A number of data quality metrics for sensor feeds have been proposed [71, 72]. Recently, a metric called *quality of contributed service* has been proposed to characterize the quality and timeliness of a sensed quantity obtained through participatory sensing [73].

- *Spatial crowdsourcing*. This is a quite new topic that proposes the engagement of users in performing certain tasks that may require them to explicitly move to specific areas to achieve the intended goals (i.e., capturing data that are linked to that geographic area) [12, 13, 14]. We can envision innovative applications that could be useful in the context of transportation, like asking vehicles to slightly detour from their route to alleviate congestion, measure some environmental parameters in an area (e.g., the CO_2), help in the task of traffic sensing, verify if a certain parking spot is still available and take a picture that can be provided to an interested car, or even just act as links in a multi-hop ad hoc data communication among vehicles. The idea is very new and should be explored in more depth.

- *Semantic data management*. The use of semantic techniques, such as ontologies [74] and reasoners [75], can provide significant benefits. Among them, we can highlight enabling the interoperability among different devices and enriching sensor data with higher-level information, which could be potentially queried in a flexible and semantic way. This is related to the so-called *Semantic Sensor Web* [76] and to concepts such as *semantic trajectory* [58] and *semantic location granule* [77]. However, how to go from raw sensor data to *smart sensor data* and exploit them effectively is an issue that requires further exploration. Ilarri et al. [61] analyze the state of the art on the semantic management of moving objects (e.g., vehicles) and propose a distributed framework for this purpose.

- *Exploitation of different types of sensors and large-scale data analytics*. How to effectively exploit and integrate all the information that sensors can provide is still an unsolved issue. For example, a geo-tagged picture posted on Facebook and local traffic data captured could be correlated to infer information about the location of the user and the traffic conditions, or a user could tweet complaining about the current service of a certain bus line in the city (which could be used to alert other people about a potential malfunctioning of a public transportation service). There is a need to identify what is relevant and how to benefit from it. The analysis of *big spatio-temporal data* to extract interesting information for urban transportation is also an avenue that deserves further exploration.

- *Different ways to cooperate, unobtrusive cooperation*. There are many different ways in which users can cooperate in the sensing process. Some of them will require explicit user interaction, but others could benefit from existing sensors in an unobtrusive way, without any effort from the user. For example, as commented in Section 3 there are several approaches to automatically detect park and unpark events that can signal the unavailability and availability of a parking space (e.g., [22, 23, 24, 25, 26]). As another example, vehicles could be used for environment monitoring without any direct action by the driver (e.g., [78, 79]). Finally, there are also several approaches for context detection, such as automatic transportation mode detection (e.g., see [59, 80]) and activity recognition (e.g., [81]), which could be exploited in collaborative transportation applications to minimize the involvement required from the user. We can also envision other intermediate possibilities that require intervention by the user but decreased to a certain extent thanks to the use of wearable devices such as the Google Glasses or the Samsung Galaxy Gear.

- *Management of trust and privacy*. Collaborative sensing for transportation can imply the collection and storage of data about users' daily trips. So, ensuring their privacy is a key issue. Similarly, trust management is a fundamental aspect of collaborative sensing in intelligent transportation, particularly if humans actively provide information. For example, malicious users could provide false information to disturb the system or to gain a competitive advantage. A survey of trust management in the transportation context including properties of trust, trust metrics, potential attacks and defenses, is presented by Ma et al. [82].

Acknowledgments

We acknowledge the support of the CICYT project TIN2013-46238-C4-4-R and DGA-FSE. This research was also supported in part by the US Department of Transportation National University Rail Center (NURAIL); and National Science Foundation grants IIS-1213013, CCF-1216096, DGE-0549489, IIP-1315169. Thanks also to the French International Campus on Safety and Intermodality in Transportation, the Nord-Pas-de-Calais Region, the European Union, the Regional Delegation for Research and Technology, the Ministry of Higher Education and Research and the National Center for Scientific Research.

References

[1] Y. Zheng, L. Capra, O. Wolfson, and H. Yang, "Urban computing: Concepts, methodologies, and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 5, no. 3, pp. 38:1–38:55, 2014.

[2] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *World Sensor Web Workshop (WSW)*. ACM, October 2006, pp. 1–5.

[3] R. K. Ganti, F. Ye, and H. Lei, "Mobile crowdsensing: Current state and future challenges," *IEEE Communications Magazine*, vol. 49, no. 11, pp. 32–39, November 2011.

[4] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *IEEE Communications Magazine*, vol. 48, no. 9, pp. 140–150, 2010.

[5] B. Fleming, "Sensors – A forecast [automotive electronics]," *IEEE Vehicular Technology Magazine*, vol. 8, no. 3, pp. 4–12, 2013.

[6] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in *Second Annual International Workshop on Wireless Internet (WICON)*. ACM, August 2006.

[7] S. Kabadayi, A. Pridgen, and C. Julien, "Virtual sensors: Abstracting data from physical sensors," in *International Symposium on on World of Wireless, Mobile and Multimedia Networks (WOWMOM)*. IEEE Computer Society, June 2006, pp. 587–592.

[8] M.-D. Albakour, C. Macdonald, and I. Ounis, "Identifying local events by using microblogs as social sensors," in *Tenth Conference on Open Research Areas in Information Retrieval (OAIR)*. ACM, May 2013, pp. 173–180.

[9] T. Sakaki, Y. Matsuo, T. Yanagihara, N. P. Chandrasiri, and K. Nawa, "Real-time event extraction for driving information from social sensors," in *International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*. IEEE Computer Society, May 2012, pp. 221–226.

[10] FP7 Project, "Socialsensor – sensing user generated input for improved media discovery and experience," `http://www.socialsensor.eu/` [Last access: April 21, 2015].

[11] M. F. Goodchild, "Citizens as sensors: The world of volunteered geography," *GeoJournal*, vol. 69, no. 4, pp. 211–221, 2007.

[12] L. Kazemi and C. Shahabi, "GeoCrowd: Enabling query answering with spatial crowdsourcing," in *20th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2012, pp. 189–198.

[13] H. To, G. Ghinita, and C. Shahabi, "A framework for protecting worker location privacy in spatial crowdsourcing," in *Proceedings of the VLDB Endowment (PVLDB)*, vol. 7, no. 10, September 2014, pp. 919–930.

[14] Z. Chen, R. Fu, Z. Zhao, Z. Liu, L. Xia, L. Chen, P. Cheng, C. C. Cao, Y. Tong, and C. J. Zhang, "gMission: A general spatial crowdsourcing platform," in *Proceedings of the VLDB Endowment (PVLDB)*, vol. 7, no. 13, September 2014, pp. 1629–1632.

[15] J. Indulska and P. Sutton, "Location management in pervasive systems," in *Australasian Information Security Workshop Conference on ACSW Frontiers (ACSW Frontiers)*, vol. 21. Australian Computer Society, Inc., January 2003, pp. 143–151.

[16] N. Islam and R. Want, "Smartphones: Past, present, and future," *IEEE Pervasive Computing*, vol. 13, no. 4, pp. 89–92, 2014.

[17] D. Shoup, "Cruising for parking," *Access*, no. 30, pp. 16–22, Spring 2007.

[18] E. Gantelet and A. Lefauconnier, "The time looking for a parking space: Strategies, associated nuisances and stakes of parking management in France," in *Europe Transport Conference (ETC)*, Association for European Transport and Contributors, 2006.

[19] M. Caliskan, D. Graupner, and M. Mauve, "Decentralized discovery of free parking places," in *Third International Workshop on Vehicular Ad Hoc Networks (VANET)*. ACM, September 2006, pp. 30–39.

[20] SFpark, "SFpark pilot project evaluation summary – a summary of the SFMTA's evaluation of the SFpark pilot project," June 2014.

[21] S. Mathur, T. Jin, N. Kasturirangan, J. Chandrasekaran, W. Xue, M. Gruteser, and W. Trappe, "ParkNet: Drive-by sensing of road-side parking statistics," in *Eighth International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, June 2010, pp. 123–136.

[22] S. Nawaz, C. Efstratiou, and C. Mascolo, "ParkSense: A smartphone based sensing system for on-street parking," in *19th Annual International Conference on Mobile Computing & Networking (MobiCom)*. ACM, June 2013, pp. 75–86.

[23] K.-C. Lan and W.-Y. Shih, "An intelligent driver location system for smart parking," *Expert Systems with Applications*, vol. 41, no. 5, pp. 2443–2456, 2014.

[24] A. Nandugudi, T. Ki, C. Nuessle, and G. Challen, "PocketParker: Pocketsourcing parking lot availability," in *International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, September 2014, pp. 963–973.

[25] B. Xu, O. Wolfson, J. Yang, L. Stenneth, P. S. Yu, and P. C. Nelson, "Real-time street parking availability estimation," in *14th International Conference on Mobile Data Management (MDM)*, vol. 1. IEEE Computer Society, June 2013, pp. 16–25.

[26] S. Ma, O. Wolfson, and B. Xu, "UPDetector: Sensing parking/unparking activities using smartphones," in *Seventh SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS)*. ACM, November 2014, pp. 1–10.

[27] B. Hoh, T. Yan, D. Ganesan, K. Tracton, T. Iwuchukwu, and J.-S. Lee, "TruCentive: A game-theoretic incentive platform for trustworthy mobile crowdsourcing parking service," in *15th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE Computer Society, September 2012, pp. 160–166.

[28] E. Kokolaki, M. Karaliopoulos, and I. Stavrakakis, "Opportunistically assisted parking service discovery: Now it helps, now it does not," *Pervasive and Mobile Computing*, vol. 8, no. 2, pp. 210–227, 2012.

[29] T. Delot, S. Ilarri, S. Lecomte, and N. Cenerario, "Sharing with caution: Managing parking spaces in vehicular networks," *Mobile Information Systems*, vol. 9, no. 1, pp. 69–98, 2013.

[30] D. Ayala, O. Wolfson, B. Xu, B. Dasgupta, and J. Lin, "Parking slot assignment games," in *19th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2011, pp. 299–308.

[31] D. Ayala, O. Wolfson, B. Xu, B. DasGupta, and J. Lin, "Pricing of parking for congestion reduction," in *20th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2012, pp. 43–51.

[32] O. Wolfson and J. Lin, "A marketplace for spatio-temporal resources and truthfulness of its users," in *Seventh SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS)*. ACM, November 2014, pp. 1–10.

[33] V. Verroios, V. Efstathiou, and A. Delis, "Reaching available public parking spaces in urban environments using ad-hoc networking," in *12th International Conference on Mobile Data Management (MDM)*. IEEE Computer Society, June 2011, pp. 141–151.

[34] A. Klappenecker, H. Lee, and J. L. Welch, "Finding available parking spaces made easy," *Ad Hoc Networks*, vol. 12, pp. 243–249, 2012.

[35] M. Caliskan, D. Graupner, and M. Mauve, "Decentralized discovery of free parking places," in *Third ACM International Workshop on Vehicular Ad Hoc Networks (VANET)*. ACM, September 2006, pp. 30–39.

[36] D. Zekri, B. Defude, and T. Delot, "Building, sharing and exploiting spatio-temporal aggregates in vehicular networks," *Mobile Information Systems*, vol. 10, no. 3, pp. 259–285, 2014.

[37] C. Morency and M. Trépanier, "Characterizing parking spaces using travel survey data," May 2008, Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), CIRRELT-2008-15.

[38] Z. Li, Y. Zhu, H. Zhu, and M. Li, "Compressive sensing approach to urban traffic sensing," in *31st International Conference on Distributed Computing Systems (ICDCS)*. IEEE Computer Society, June 2011, pp. 889–898.

[39] A. Thiagarajan, L. Ravindranath, K. LaCurts, S. Madden, H. Balakrishnan, S. Toledo, and J. Eriksson, "VTrack: Accurate, energy-aware road traffic delay estimation using mobile phones," in *Seventh ACM Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, November 2009, pp. 85–98.

[40] M. Tanizaki and O. Wolfson, "Randomization in traffic information sharing systems," in *15th SIGSPATIAL International Symposium on Advances in Geographic Information Systems (GIS)*. ACM, November 2007, pp. 23:1–23:8.

[41] D. Ayala, J. Lin, O. Wolfson, N. Rishe, and M. Tanizaki, "Communication reduction for floating car data-based traffic information systems," in *Second International Conference on Advanced Geographic Information Systems, Applications, and Services (GEOPROCESSING)*, February 2010, pp. 44–51.

[42] R. Du, C. Chen, B. Yang, N. Lu, X. Guan, and X. Shen, "Effective urban traffic monitoring by vehicular sensor networks," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 1, pp. 273–286, January 2015.

[43] J. A. Sanguesa, M. Fogue, P. Garrido, F. J. Martinez, J.-C. Cano, C. T. Calafate, and P. Manzoni, "An infrastructureless approach to estimate vehicular density in urban environments," *Sensors*, vol. 13, no. 2, pp. 2399–2418, 2013.

[44] J. Barrachina, J. A. Sanguesa, M. Fogue, P. Garrido, F. J. Martinez, J.-C. Cano, C. T. Calafate, and P. Manzoni, "V2X-d: A vehicular density estimation system that combines V2V and V2I communications," in *IFIP Wireless Days (WD)*, November 2013, pp. 1–6.

[45] J.-Y. Yang, L.-D. Chou, C.-F. Tung, S.-M. Huang, and T.-W. Wang, "Average-speed forecast and adjustment via VANETs," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 9, pp. 4318–4327, 2013.

[46] S. Djahel, R. Doolan, G.-M. Muntean, and J. Murphy, "A communications-oriented perspective on traffic management systems for smart cities: Challenges and innovative approaches," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 125–151, 2015.

[47] B. Pan, Y. Zheng, D. Wilkie, and C. Shahabi, "Crowd sensing of traffic anomalies based on human mobility and social media," in *21st SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2013, pp. 344–353.

[48] R. Varriale, S. Ma, and O. Wolfson, "VTIS: A volunteered travelers information system," in *Sixth SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS)*. ACM, November 2013, pp. 13:13–13:18.

[49] H. van Lint, "Reliable travel time prediction for freeways – bridging artificial neural networks and traffic flow theory," Ph.D. dissertation, Delft University of Technology, Delft (Netherlands), May 2004.

[50] S. M. Bilal, C. J. Bernardos, and C. Guerrero, "Position-based routing in vehicular networks: A survey," *Journal of Network and Computer Applications*, vol. 36, no. 2, pp. 685–697, 2013.

[51] C. S. Chong, B. Zoebir, A. Y. S. Tan, W.-C. Tjhi, T. Zhang, K. K. Lee, R. M. Li, W. L. Tung, and F. B.-S. Lee, "Collaborative analytics for predicting expressway-traffic congestion," in *14th Annual International Conference on Electronic Commerce (ICEC)*. ACM, August 2012, pp. 35–38.

[52] B. Xu and O. Wolfson, "Time-series prediction with applications to traffic and moving objects databases," in *Third ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDe)*. ACM, September 2003, pp. 56–60.

[53] W. Min and L. Wynter, "Real-time road traffic prediction with spatio-temporal correlations," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 4, pp. 606–616, 2011.

[54] O. Wolfson and B. Xu, "A new paradigm for querying blobs in vehicular networks," *IEEE MultiMedia*, vol. 21, no. 1, pp. 48–58, 2014.

[55] G. Sun, J. Chen, W. Guo, and K. Liu, "Signal processing techniques in network-aided positioning: A survey of state-of-the-art positioning designs," *IEEE Signal Processing Magazine*, vol. 22, no. 4, pp. 12–23, 2005.

[56] X. Lu, C. Wang, J.-M. Yang, Y. Pang, and L. Zhang, "Photo2Trip: Generating travel routes from geotagged photos for trip planning," in *International Conference on Multimedia (MM)*. ACM, October 2010, pp. 143–152.

[57] H. Yin and O. Wolfson, "A weight-based map matching method in moving objects databases," in *16th International Conference on Scientific and Statistical Database Management (SSDBM)*. IEEE Computer Society, June 2004, pp. 437–438.

[58] C. Parent, S. Spaccapietra, C. Renso, G. Andrienko, N. Andrienko, V. Bogorny, M. L. Damiani, A. Gkoulalas-Divanis, J. Macedo, N. Pelekis, Y. Theodoridis, and Z. Yan, "Semantic trajectories modeling and analysis," *ACM Computing Surveys*, vol. 45, no. 4, pp. 42:1–42:32, 2013.

[59] L. Stenneth, O. Wolfson, P. S. Yu, and B. Xu, "Transportation mode detection using mobile phones and GIS information," in *19th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2011, pp. 54–63.

[60] J. Liu, O. Wolfson, and H. Yin, "Extracting semantic location from outdoor positioning systems," in *Seventh International Conference on Mobile Data Management (MDM)*, May 2006, pp. 73–73.

[61] S. Ilarri, D. Stojanovic, and C. Ray, "Semantic management of moving objects: A vision towards smart mobility," *Expert Systems With Applications*, vol. 42, no. 3, pp. 1418–1435, 2015.

[62] F. Xu, S. Guo, J. Jeong, Y. Gu, Q. Cao, M. Liu, and T. He, "Utilizing shared vehicle trajectories for data forwarding in vehicular networks," in *30th International Conference on Computer Communications (INFOCOM)*. IEEE Computer Society, April 2011, pp. 441–445.

[63] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang, "T-Drive: Driving directions based on taxi trajectories," in *18th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*. ACM, November 2010, pp. 99–108.

[64] J. J.-C. Ying, E. H.-C. Lu, W.-C. Lee, T.-C. Weng, and V. S. Tseng, "Mining user similarity from semantic trajectories," in *Second SIGSPATIAL International Workshop on Location Based Social Networks (LBSN)*. ACM, November 2010, pp. 19–26.

[65] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, "Mining interesting locations and travel sequences from GPS trajectories," in *18th International Conference on World Wide Web (WWW)*. ACM, April 2009, pp. 791–800.

[66] W. He, K. Hwang, and D. Li, "Intelligent carpool routing for urban ridesharing by mining GPS trajectories," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 5, pp. 2286–2296, 2014.

[67] S. Ma, Y. Zheng, and O. Wolfson, "T-Share: A large-scale dynamic taxi ridesharing service," in *29th International Conference on Data Engineering (ICDE)*, April 2013, pp. 410–421.

[68] F. Giannotti, M. Nanni, D. Pedreschi, F. Pinelli, C. Renso, S. Rinzivillo, and R. Trasarti, "Unveiling the complexity of human mobility by querying and mining massive trajectory data," *The VLDB Journal*, vol. 20, no. 5, pp. 695–719, October 2011.

[69] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing," in *18th Annual International Conference on Mobile Computing and Networking (Mobicom)*. ACM, August 2012, pp. 173–184.

[70] K. Sugiyama, T. Hasegawa, J. Huang, T. Kubo, and J. Walrand, "Motivating smartphone collaboration in data acquisition and distributed computing," *IEEE Transactions on Mobile Computing*, vol. 13, no. 10, pp. 2320–2333, 2014.

[71] L. Ramaswamy, V. Lawson, and S. V. Gogineni, "Towards a quality-centric Big Data architecture for federated sensor services," in *2013 IEEE International Congress on Big Data (BigData Congress)*. IEEE Computer Society, June 2013, pp. 86–93.

[72] Z. Qin, Q. Han, S. Mehrotra, and N. Venkatasubramanian, “Quality-aware sensor data management,” in *The Art of Wireless Sensor Networks*, ser. Signals and Communication Technology, H. M. Ammari, Ed. Springer, 2014, pp. 429–464.

[73] C. Tham and T. Luo, “Quality of contributed service and market equilibrium for participatory sensing,” *IEEE Transactions on Mobile Computing*, vol. 14, no. 4, pp. 829–842, 2015.

[74] I. Horrocks, “Ontologies and the Semantic Web,” *Communications of the ACM*, vol. 51, no. 12, pp. 58–67, 2008.

[75] R. B. Mishra and S. Kumar, “Semantic web reasoners and languages,” *Artificial Intelligence Review*, vol. 35, no. 4, pp. 339–368, 2011.

[76] J.-P. Calbimonte, H. Jeung, Ó. Corcho, and K. Aberer, “Enabling query technologies for the Semantic Sensor Web,” *Int. Journal on Semantic Web and Information Systems*, vol. 8, no. 1, pp. 43–63, 2012.

[77] J. Bernad, C. Bobed, E. Mena, and S. Ilarri, “A formalization for semantic location granules,” *Int. Journal of Geographical Information Science*, vol. 27, no. 6, pp. 1090–1108, 2013.

[78] O. Urra, S. Ilarri, E. Mena, and T. Delot, “Using hitchhiker mobile agents for environment monitoring,” in *Seventh International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*, ser. Advances in Intelligent and Soft Computing, vol. 55. Springer, March 2009, pp. 557–566.

[79] X. Xu, P. Zhang, and L. Zhang, “Gotcha: A mobile urban sensing system,” in *12th ACM Conference on Embedded Network Sensor Systems (SenSys)*. ACM, November 2014, pp. 316–317.

[80] S. Reddy, M. Mun, J. Burke, D. Estrin, M. Hansen, and M. Srivastava, “Using mobile phones to determine transportation modes,” *ACM Transactions on Sensor Networks*, vol. 6, no. 2, pp. 13:1–13:27, 2010.

[81] N. Győrbíró, Á. Fábián, and G. Hományi, “An activity recognition system for mobile phones,” *Mobile Networks and Applications*, vol. 14, no. 1, pp. 82–91, 2008.

[82] S. Ma, O. Wolfson, and J. Lin, “A survey on trust management for intelligent transportation system,” in *Fourth SIGSPATIAL International Workshop on Computational Transportation Science (CTS)*. ACM, November 2011, pp. 18–23.