Crowdsensing Smart Ambient Environments and Services

Blake Regalia, Grant McKenzie, Song Gao, and Krzysztof Janowicz

STKO Lab, Department of Geography,
University of California, Santa Barbara, CA, USA

Abstract

Whether it be Smart Cities, Ambient Intelligence, or the Internet of Things, current visions for future urban spaces share a common core, namely the increasing role of distributed sensor networks and the on-demand integration of their data to power real-time services and analytics. Some of the greatest hurdles to implementing these visions include security risks, user privacy, scalability, the integration of heterogeneous data, and financial cost. In this work, we propose a crowdsensing mobile-device platform that empowers citizens to collect and share information about their surrounding environment via embedded sensor technologies. This approach allows a variety of urban areas (e.g., university campuses, shopping malls, city centers, suburbs) to become equipped with a free ad-hoc sensor network without depending on proprietary instrumentation. We present a framework, namely the GeoTracer application, as a proof of concept to conduct multiple experiments simulating use-case scenarios on a university campus. First, we demonstrate that ambient sensors (e.g., temperature, pressure, humidity, magnetism, illuminance, and audio) can help determine a change in environment (e.g., moving from indoors to outdoors, or floor changes inside buildings) more accurately than typical positioning technologies (e.g., global navigation satellite system, Wi-Fi, etc.). Furthermore, each of these sensors contributes a different amount of data to detecting events. for example, illuminance has the highest *information gain* when trying to detect changes between indoors and outdoors. Second, we show that through this platform it is possible to detect and differentiate place types on a university campus based on inferences made through ambient sensors. Lastly, we train classifiers to determine the activities that a place can *afford* at different times (e.g., good for studying or not, basketball courts in use or empty) based on sensor-driven *semantic signatures*.

Keywords: crowdsensing, smart city, semantics, volunteered geographic services

1   Introduction

While the vision of smart homes and cities is not new (Edwards and Grinter, 2001; Komninos, 2002), the core theme of unifying sensor networks, cyber-infrastructures, interoperability, and predictive analytics research for the purpose of improving the quality of life in urban systems is rapidly gaining momentum. As technology has changed, so has the concept behind smart environments. A

growth in sensor capabilities, wireless communication, and battery efficiency, combined with the increasing availability of (personal) data and methods for their synthesis lead to a wide adoption of these technologies for the purpose of enhancing the urban environment (Ganti et al., 2011). As smart cities and homes are taking shape, there is increased interest from the research community, industry, and governments regarding the technical challenges, social implications, and practical applications brought on by sensors and their interlinked observations (Martins and McCann, 2015).

One of the primary barriers to entry facing cities, university campuses, home owners, and organizations looking to adopt intelligent sensors is the instrumentation infrastructure. The deployment and maintenance of large-scale, interconnected sensor networks come with significant financial overhead. Not only are most authorities required to notify the public of any pending changes to public space, but they are also required to entertain bids from a variety of third party companies. Installation of sensors and transmitters can cause years of disruption to the urban ecosystem both physically, through construction, and politically, through various levels of *red tape*. As more cities adopt intelligent sensor platforms, issues of security have also begun to emerge. For instance, smart cities are vulnerable to cyber-attacks at different levels (Elmaghraby and Losavio, 2014). The very premise of a smart city dictates that everything must be connected to everything else; which can be maliciously exploited to aid criminal activity ranging from identify theft (Sen et al., 2013) to stop light resequencing (Cerrudo, 2015).

The adoption of sensor technologies by governments and the industry has also raised concerns over privacy issues. Ubiquitous technologies such as video surveillance and traffic sensors constantly monitor human activities in urban settings. The intended purpose of these systems might be to aid emergency responders or to improve traffic management but the technologies also enable nefarious activities. To give a concrete and representative example, consider the *smart-bin* technology. Prior to the 2012 London Olympics, new garbage bins were installed throughout the city to improve the availability of depositing trash. The bins, however, also housed an LCD screen that was used as advertising space. It was eventually discovered that these bins were also collecting MAC addresses from the mobile devices of citizens as they walked around the city. Over 4 million device captures took place without the knowledge nor expressed permission of device owners or city officials (Vincent, 2013). This gave a small group of people the ability to track hundreds of thousands of citizens throughout the city and model their daily activities. In many ways, citizens of today's smart cities have paid for the benefits of amenities like reduced traffic congestion at the cost of their personal privacy. Other potential problems arising from the top-down deployment of sensor networks and their services include issues of access, payment, and prioritization to said services and derived products, e.g., user traces.

It has become clear that a positive definition of the smart city needs to include citizen engagement at all horizontal and vertical levels, placing citizens at the focal point of the city. It is here that we take up Michael Goodchild's original call to *citizens as sensors* (2007). His work introduced the concept of Volunteered Geographic Information (VGI) as a way for individuals to contribute to their environment through user-contributed geographic content. Our work approaches the citizens as sensors idea from a more literal perspective, pointing towards the sensors accessible via an individual's mobile device, e.g., a smartphone, smartwatch, tablet, or smart clothing. Many applications have been developed to take advantage of these sensors ranging from hyper-sensitive mobile gaming platforms that make use of accelerometers to barometric pressure sensors that report elevation change to hiking apps. Until now however, the vast majority of these sensor-centric applications have targeted the individual rather than society as a whole, a notable exception be-

ing crowdsensing systems such as McSense (Cardone et al., 2013). In this work, we argue that sensor-rich devices have lowered the barrier to entry for the democratization of smart cities and other environments more broadly. Citizens now have the ability and opportunity to collect data via their existing devices and can share the information with all partnering citizens, enriching the environment in which they have a vested, joint interest.

Citizen engagement allows for places to be redefined by how they are used in practice rather than how city planners intended their use. This has far reaching implications, e.g., on the ways how we conceptualize and organize places by types, e.g., in geo-ontologies, and therefore what kind of queries we can pose to our intelligent personal assistants (Chaudhri et al., 2006). A better understanding of affordance-based place models (Jordan et al., 1998; Scheider and Janowicz, 2014) can be ascertained through the sensors carried by the people who occupy the spaces. For example, a city park may have been originally designed for use on weekends, but data collected by visitors indicates that the park is quite popular for after school programs. A cafeteria that affords eating during lunch hours might be ideal for studying in the early mornings and late evenings. In principle, any public location indoor or outdoor can be used for studying as long as it fits the profile, e.g. sufficiently bright and quiet — a combination of conditions from which a smartphone's luminosity sensor and microphone can easily produce a digital signature.

An added benefit of including citizens as sensors in the smart city system is that the quantity of sensor contributions for a particular area can reflect the usage statistics of public spaces. As will be discussed throughout this paper, we envision an ad-hoc network of citizens contributing sensor data as they go about their daily lives. In essence, these citizens are creating ambient smart environments around the places and locations that are important to them. In other words, every environment ranging from a university campus to a campground can be turned into a smart ambient environment without the need for top-down infrastructure and instrumentation. This is also of importance in case of extreme events such as earthquakes.

In the interest of protecting privacy, the proposed ad-hoc sensor-network does not require human identification. A second's worth of anonymized sensor readings from a limited set of distributed devices would be enough to build a rich, hyper-local (Xia et al., 2014) model of an ambient field such as illuminance. Crowdsourced sensor readings compose the backbone of a community-driven service-oriented platform, a concept envisioned by Volunteered Geographic Services (Savelyev et al., 2011) and other systems before.

Summing up, today's approaches towards smart cities and smart environments more generally follow a static top-down approach in which decisions are made well in advance and based on criteria such as the availability of resources. Once the physical infrastructure, e.g., sensor networks, have been deployed, the abilities and services offered by the smart city become fixed and immutable, oftentimes controlled solely by the governments and instrument vendors which raises well justified concerns among citizens.

2   Research Contribution

In this work, we propose a true citizens-as-sensors platform that allows individuals to contribute information collected via mobile sensors for the purpose of generating hyper-local smart ambient environments on-the-fly by deploying services on top of signatures learned from the jointly collected sensor observations. Thereby we combine ongoing research on learning place types via

bottom-up approaches, crowdsensing, VGI, and so forth. To give a first intuition, instead of relying on a handful of temperature sensors in a city, we envision a platform in which citizens can open an application to sample the surrounding environment passively and anonymously as they move throughout their locality and combine these data across sensor types and citizens.

Today's entry-level smartphone includes an impressive array of reasonably accurate sensors capable of recording from several different ambient fields. We have reached a point in mobile device development where we can now think more intelligently about how we might use these sensors to enhance our everyday lives, e.g., by detecting and reacting to activities. Battery expenditure is still very much an issue for the majority of consumer-grade mobile devices and recent research has shown that many of the low-battery usage sensors can be used to enhance and sometimes replace higher-battery usage sensors. For example, GPS receivers are among the most energy-expensive sensors on a typical mobile device. Previous work (Zandbergen, 2009) has shown that Wi-Fi positioning and Cellular trilateration alone do not yield reasonably precise location fixes, however using a combination of methods (Lin et al., 2010) can improve positioning. As we will demonstrate in this work, sensor readings of temperature, humidity, and illuminance can further reduce the uncertainty associated with positioning and also differentiate between adjacent spaces that may otherwise confuse typical positioning technologies.

The research contribution of this work is rooted in the concept of *place*. While sensor readings are associated with spatial fixes, it is through the *activities* of individuals at those locations that the true power of signatures are exposed. The integration of observations from different sensor types yields unique semantic signatures that can be used to differentiate places based on the types of activities that they afford. Even nuanced differences between sensor readings (e.g., barometric pressure) often indicate changes in environmental conditions such as moving from one floor in a building to another. Similarly, when positioning technologies locate an individual inside of a shopping mall, signatures generated from audio sensors alone provide the ability to distinguish between say a restaurant and a book store. As we examine the various combinations of sensors and the ways in which they contribute to place signatures, it is important to explore the degree of each sensor's contribution to these signatures that allows us to differentiate place types. While using all available sensors in concert would certainly improve the quality and robustness of place type estimations, it would come at the costs of increased battery consumption and system resource usage.

Summing up, we propose a hyper-local, sensor-driven platform that uses data from the sensors of mobile devices to construct ambient smart environments and offers services enabled by the signatures collected by these devices.

The novel contributions of this research are as follows:

1. *Enhance existing location-aware technologies.*
   We show that ambient sensors accessible on most mobile devices can determine a change in environment (e.g., moving from indoors to outdoors, or changing floors inside a building) more accurately than standard location technologies (e.g., GNSS, WiFi).

2. *Detect and differentiate place types*
   We show that through this platform it is possible to detect and differentiate place types on a university campus based purely on sensors readings from a mobile device.

3. *Enable hyper-local activity services*
   We show that the proposed platform permits user-generated geographic services (Savelyev et al., 2011) that rely on up-to-date information for enhancing place-based activities. For instance, one can trigger distant sensors to detect whether a class is already in session or not using the sensors of mobile devices and machine learning. This approach fits into the broader vision of constructing hyper-local environmental fields as well as services for querying this information.

The remainder of this paper is organized as follows. In Section 3 we discuss existing work related to our proposed research. A discussion of mobile device sensors and the methods used to extract information from these sensors is introduced in Section 4. An overview of the mobile platform designed to collect and assess these sensors readings is described in Section 5. Example scenarios for this framework are outlined and explored in Section 6 and conclusions and future work are presented in Section 7.

3   Related Work

Todays smartphones are equipped with a growing set of powerful embedded sensors which advance the emerging field of mobile phone sensing and facilitate several services such as social networking, environmental monitoring and health care (Lane et al., 2010). Location-based services also benefit from embedded geographic and proximity-based sensors. Research has also shown that accelerometer data can be used to recognize different activity modes (e.g., walking, running, biking) (Wang et al., 2009a), which has been integrated with several innovative fitness-related technologies and applications.

Microphones also offer a unique perspective on the context of place, which have been used to measure noise pollution levels (Santini et al., 2009) and to monitor road conditions by analyzing recorded audio snippets, such as Microsoft's *Nericell* project (Mohan et al., 2008). In addition, some researchers have demonstrated that by using ubiquitous sensing data collected from human sensors, they can build models to diagnose urban noise. For example, Zheng et al. (2014) propose a three-dimensional (consisting of regions, noise categories, and time slots) tensor model to infer urban noise conditions at different times of the day for several regions in New York City, by using 3-1-1 compliant data combined with social media, road network data, and points of interests (POIs).

Savelyev et al. (2011) coined a new concept of volunteered geographic services (VGS) in which users contribute to location-based content while simultaneously improving crowdsourced service-oriented platforms. They develop a novel architecture using the Linked Data paradigm and Semantic Web technologies to support location-based human-oriented services, such as ridesharing.

McKenzie and Janowicz (2015) develop a location-distortion model to improve reverse geocoding technology using behavior-driven temporal semantic signatures extracted from online location-based social networks. Either the combination or one of spaital, temporal and thematic bands have been demonstrated to play an important role in differentiating place types (McKenzie et al., 2015).

Positioning is the key component to support location-based studies. It is well known that GPS-enabled devices work well for positioning in outdoor environments but not in indoor environments because of the signal obstructions. However, as reported in a national human activity pattern survey

(NHAPS), people spend an average of 87% of their time in enclosed buildings and about 6% of their time in enclosed vehicles (Klepeis et al., 2001).

Many researchers have made great efforts in the field of indoor positioning. Elhamshary and Youssef (2015) develop a crowdsourcing-based system that can automate the construction of semantic-rich indoor floorplans with a specific label (e.g., venue name in a mall) or a general place category (e.g., restaurant).

Mazilu et al. (2013) propose to use low-power ambient sensors: temperature (T), humidity (H), pressure (P), and light (L) integrated in phones to extract semantic locations (such as home, office, train, and shop), which can be an alternative to standard power-hungry localization methods, such as GPS/WiFi/GSM and audio, while significantly reducing the power consumption and saving up to 65% of battery power for the THP sensor combination, and up to 85% for the TH combination.

Cardone et al. (2013) introduce a framework to collect sensor data from a distributed sensor-network of mobile devices in which participants are assigned tasks for the purpose of enabling data-driven geo-services. We describe here a similar product that extends the conceptual device platform with ambient sensors for passive data collection and learns place type signatures bottom-up to support various location-based services.

4   Methodology

In this section we introduce sensors that are common to many mobile devices and discuss a number of approaches that can be used to extract signatures from these sensors. We then show that through a variety of statistical and machine learning methods, these signatures can be employed to identify a range of human activities and place types.

4.1   Device Sensors

Most current mobile devices have built-in sensors that measure location (GNSS), motion (e.g., accelerometer, gyroscope), radio-frequency signals (e.g., Bluetooth, Wi-Fi), and various ambient environmental conditions (e.g., microphone, light sensor, barometer, hydrometer). These sensors are capable of providing raw timestamped data to monitor device orientation, position, and changes to the surrounding environment. The dominant mobile operating system world-wide as of 2015[1] is Android. By examining publicly-available statistics about the Android operation-system version distributions[2], we can identify prevalence among certain sensors. In rare cases, new releases of Android come with an explicit recommendation to hardware manufacturers to drop a sensor from their device, such as the *Compatibility Definition* document of Android 4.2 that designated the temperature sensor be deprecated [3]. A recent report found that the most prevalent sensor on Android mobile devices world-wide was the accelerometer with 30% adoption having increased adoption substantially from 2012. This was followed by barometric pressure sensor which was housed by 20% of Android devices followed by magnetic field sensors and humidity sensors OpenSignal

---

[1]https://www.idc.com/prodserv/smartphone-os-market-share.jsp

[2]http://developer.android.com/intl/ru/about/dashboards/index.html

[3]http://static.googleusercontent.com/media/source.android.com/en/
/compatibility/4.2/android-4.2-cdd.pdf

(2015). Note that not all sensors were mentioned in this report and that these low percentages include older devices running very early versions of the Android operating system or systems other than smartphones and tablets.

Mobile device sensors have different power consumption rates. GNSS for example has been shown to consume approximately ten times more power than Wi-Fi under idle conditions (Carroll and Heiser, 2010). Knowing this, many studies have made efforts to develop energy efficient mobile sensing platforms for positioning, tracking or automatic user state detection (Wang et al., 2009b; Paek et al., 2010; Bhattacharya et al., 2015). Compared to energy-expensive sensors like GNSS, sensors such as accelerometers, gyroscopes, barometric altimeters, and microphones exhibit relatively low power-consumption behavior, though the power consumption of these sensors also varies considerably by device. For example, it has been reported that the actual power cost to continuously gather data from the Android barometer is about 1000uW while 5000uW is needed for the gyroscope[4]. More importantly, it has been shown that components involved with radio transmitting and receiving (e.g., WiFi, Bluetooth, GSM, GNSS, etc.) are the predominant power-consumers among all other sensor components for typical applications (Carroll and Heiser, 2010).

The availability of sensors is an important part of the ambient sensor platform but the power consumption should also be considered when used in practice. In some cases, it is much more efficient to use a combination of low energy sensors in lieu of a single more energy-expensive one. As will be shown in the next section (with regards to GNSS), this also leads to more accurate results.

4.2 Signatures and Bands

Different types of sensors generate varying data in different environment settings. In previous work, we proposed the concept of *semantic signatures* to characterize *types* of places using spatial, temporal, and thematic bands (Janowicz, 2012). We refer to such multi-dimensional quantitative information as *semantic signatures*. As an analogy to *spectral signatures* in remote sensing, *semantic signatures* differentiate types of places based on multiple bands. In this work, we take those unique signatures from different types of mobile sensors as semantic bands to differentiate place types and environment settings. To give an intuitive example, a temporal band (e.g., hourly check-ins of users to places of certain types) can be used to readily distinguish nightclubs from restaurants but not necessarily nightclubs from bars. By combining additional bands, however, such as thematic bands that represent how people describe places (of given types), such distinctions become possible. Hence, a signature is just the combination of bands that uniquely identifies a type of geographic feature. In the work at hand, the bands of the signatures are not check-in times or textual descriptions but readings from multiple sensor types.

4.3 Statistical Approach

Approaching a set of mobile device sensor readings as an array of signatures, machine learning is used to both extract patterns in the data as well as build predictive models for activity and place (type) identification. These methods are further described through examples in Section 6.

---

[4]https://source.android.com/devices/sensors/sensor-types.html

4.3.1 Entropy and Information Gain

We use an information entropy-based approach for determining the amount of information that each sensor contributed to the classification of place types. In our case, *Shannon information entropy* (Equation 1) is used to quantify the amount of variability in a given sensor's readings (Shannon, 1951).

$$H(D) = -\sum_{i=1}^{n} p_i log_2(p_i) \tag{1}$$

$H(D)$ in this equation represents the entropy of a given sensor reading while p_i represents the probability of a particular value belonging to a class *i*.

Building on the knowledge that each sensor contributes a certain amount to a model used to differentiate places (for example), we then calculate the *information gain* (IG) (Quinlan, 1986) for each sensor. This approach in essence reports the relative amount by which each sensor contributes to the classification, e.g., the indoor or outdoor status. Equation 2 shows the information gain formula for a specific sensor j with $\frac{|D_j|}{|D|}$ being the weight of the *jth* sensor of the training set.

$$IG(D|D_j) = H(D) - \sum_{j=v}^{n} \frac{|D_j|}{|D|} \times H(D_j) \tag{2}$$

4.3.2 Support Vector Machine

In this paper, we will demonstrate how leveraging mobile sensor information from a crowdsensing network helps determine suitable locations to conduct certain types of activities. For cases such as this, it is appropriate to build a classification model using sample data for training. A multi-class support vector machine (SVM) (Suykens and Vandewalle, 1999; Kuhn, 2008) is a type of supervised learning model that we use by feeding sensor readings from labeled training data as inputs to build a classification model which can then predict and assign labels to untagged sensor readings. In Section 6 this approach is used to assess the suitability of a place (type) for *studying* provided sensor readings from known study locations.

4.4 Organizing the Implications of Context by Engineering an Ontology

The platform that we propose in this work collects sensor readings from a network of volunteers who annotate their data by tagging each recording with text labels that identify enviornmental qualities (e.g., quiet, crowded, dark, etc.) as well as the current place or activity types (e.g., studying, eating, library, cafeteria, etc.). Although we train our model to make accurate predictions about activity types given ambient sensor signatures, the device is still far removed from the implications of such category 'tags'. In order to build geographic services that provide practical features to the users, we must be able to model how activities affect place, which actions should be triggered by certain conditions, and how to appreciate the full hierarchy of context.

When it comes to organizing the high-level relationships that associate sensor reading qualities to the affordances of place types, an ontology is a common way to structure such information. While traditional methods may take an exclusive top-down approach to the engineering of such ontologies, we propose to engineer them also bottom-up, i.e., by mining semantic relations from
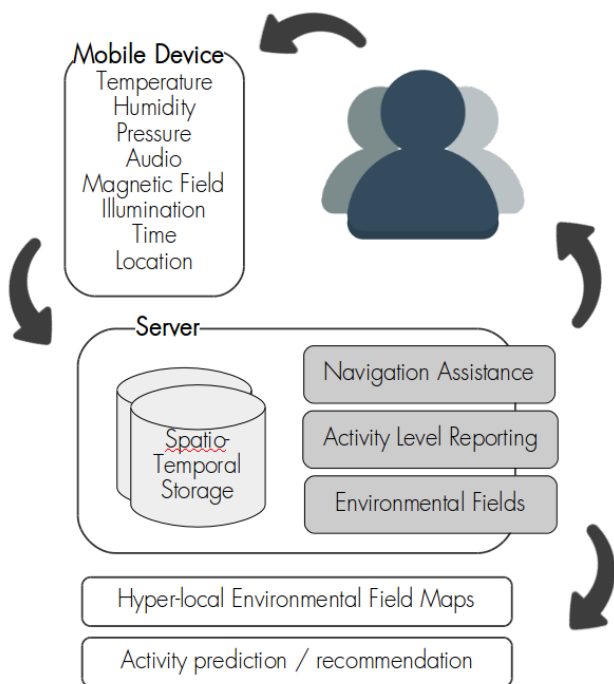
Figure 1: Overview of the proposed ambient sensor framework.

massive amounts of user-contributed sensor data. One such strategy is to start with hierarchical clustering on the labels, where the distance between elements is given by their co-occurrence among sensor readings and then derive similarity relations between clusters.

A simple example to consider is a mobile device in a movie theater. With sufficient context, a proper ontology can indicate that an appropriate action would be to enter silent mode. The ontology itself does not need an explicit association between movie theater and silent mode; rather, *movie theater* may hold association to the class of *respectfully quiet place*, which in turn points to the action of *enter silent mode*. Rather than hard coding these rules into software, one can organize the implications of activities and context in an ontology that is equally top-down (theory-driven) and bottom-up (data-driven); see (Keßler et al., 2009) for an example.

5 Implementation: The *GeoTracer* Platform

For this research project, we developed an open source data collection application for Android because its platform supports many of the ambient sensors we plan to use. The application is written in Java to make use of Android's application programming interface (API) for collecting data from sensor hardware. Since we plan to reuse our code for the final product, it is designed to run with a minimal system resource footprint including: low power consumption, efficient memory usage, and a data encoding format that spares storage space (714kb file for 67.5k data events for a 30 minute period of continuous sampling). We also implemented a complementary decoder in Java which runs on a centralized server where the recorded data files get sent to by each device. The decoder is capable of outputting to several different formats including SQL, JSON, and CSV.

Since the data collection application needs to record sensor readings from indoor locations, we

had to associate each sample with the place from which it was collected. Rather than building an indoor positioning system to locate each sample spatially, we used floor plans of the focal buildings allowing us to manually specify our location to the application's interface. Since our objective was to identify place types and detect changes in activity, the spatial accuracy was less important. Along with this interface, we also included a mechanism that allowed us to tag the types of activities that occurred at that place and time to establish a training dataset. Toggling activity labels during our trials allowed us to associate sensor signatures with an activity. We then used tagged activity data to train a place type classification model. Using the data we collected in the field, our model then computes the alignment of each sensor signature to its corresponding classification of place type, activity, and so on (see Figure 1). The output from this model is what enables a *smart environment* to make predictions about the activities afforded by various places at given times, the primary feature of such a framework. This prototype application is the first iteration of the proposed platform. The version of the application currently in development allows authenticated users to annotate their sensor readings by toggling labeled tags as specific places based on their activity type. This improves the place type signatures over time as more data is collected from a diverse variety of environments and thus makes the classification more robust.

6 Scenarios

The methods and framework discussed in the previous sections present the value of an ambient sensor-driven framework and describe our approach at an abstract level. In this section we introduce a number of *scenarios* that demonstrate the value of the information accessible via mobile ambient sensors. Three scenarios are described along with some real-world sensor data examples against which some use cases are validated. The purpose of these example scenarios is to give a representative overview as to what is possible through such a platform.

6.1 Enhancing Location Technologies

Current approaches to navigation and place detection rely heavily on the use of radio technologies such as GNSS, cellular trilateration (e.g., GSM, CDMA), Wi-Fi or Bluetooth. Each of these approaches, however, comes with its own strengths and weakness. For example, GNSS requires line-of-sight between the receiver and multiple satellites, making it inadequate for indoor positioning. Wi-Fi and Bluetooth methods require an ample number of beacons in an ideal spatial configuration to cover a service area, and are susceptible to marginal positioning errors such as locating a user on the wrong floor in a building. We propose to use ambient sensors on a mobile device to complement current location technologies and in many cases, enhance them.

6.1.1 Scenario: Indoor vs. Outdoor

The first scenario that we investigate for enhancing location technologies is detecting place change events that might challenge the credibility of location providing services, such as the boundary between moving from outdoors to indoors where GNSS signals can no longer be trusted. Our argument follows that the mere indication of such a place change event is reason enough to switch off the GNSS sensor, preserve energy, and prevent location-consuming applications from being

misled about the user's trajectory. Existing work has shown that ambient sensors on mobile devices can reliably be used to detect room changes (Mazilu et al., 2013), for example walking from a kitchen into a bedroom. Following this direction, and to support our assumption about GNSS reliability, we demonstrate that the same effect holds for detecting a place change event from outdoors to indoors.

Figure 2 shows two real-world examples scenarios. In the first case (red dots), a user with a GPS-enabled mobile device walks past a building. In the second case (blue dots) the user walks into the building. These dots are the *reported* GPS fixes from these two scenarios. The solid line data shows the user's *actual* trajectories. This figure shows that provided GPS-based trajectories alone, it would be difficult to determine whether or not the user entered or walked past the building.
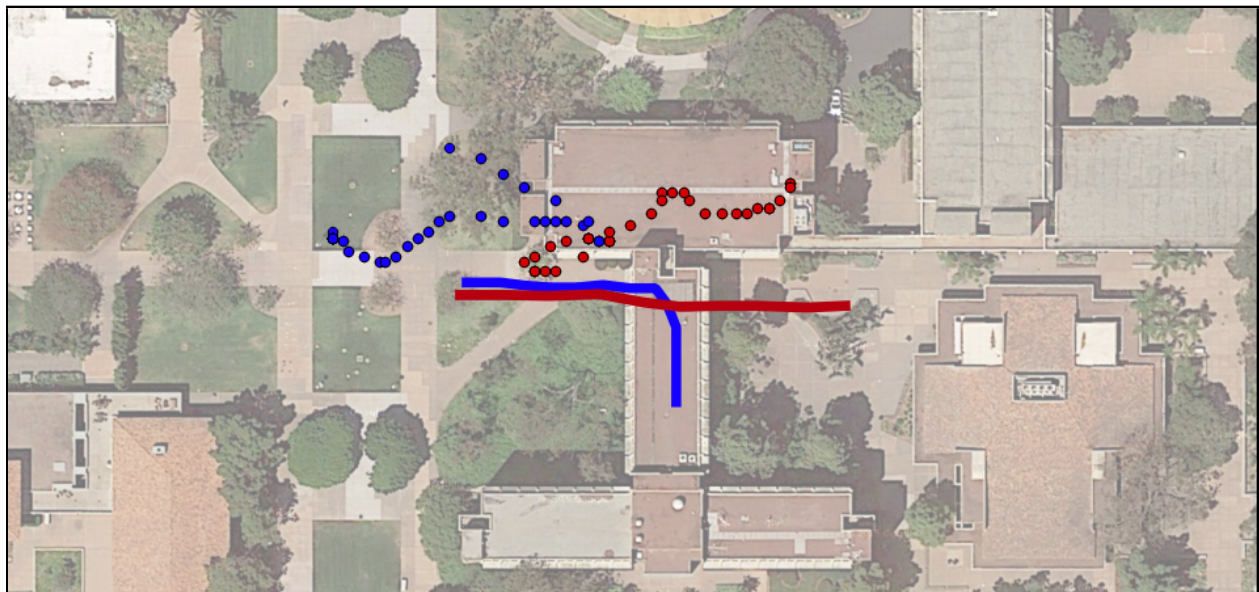

Figure 2: GPS coordinates for two scenarios. The first (red dots) shows someone walking past a building and the second (red dots) shows someone walking into the building. The solid lines represent the actual trajectories.

Ambient field sensors on a user's mobile device, however, can be exploited to differentiate these two activities. As introduced above, a sample trajectory was collected from a user walking outside of a building into the indoor hallway. Sensor data was recorded at 500 millisecond intervals and each sensor reading was tagged as either being inside or outside of the building. The switch from outside to inside was made at the moment that the user crossed the plane of the building's doorway. Figure 3 visually depicts the sensor information reported over the course of these test runs. As one might expect, there are some clear differences between indoor and outdoor environments that are reflected in the sensor readings. *Light* is an important indicator in detecting this change. During the moment a user crosses the indoor/outdoor boundary, the luminous flux per unit area (lux) changes significantly due to a change from direct sunlight to indoor florescent lights. Additionally, the temperature change between the two environments is significant enough to indicate a place change event.

The results of an information gain approach are shown in Table 1. The data in this table
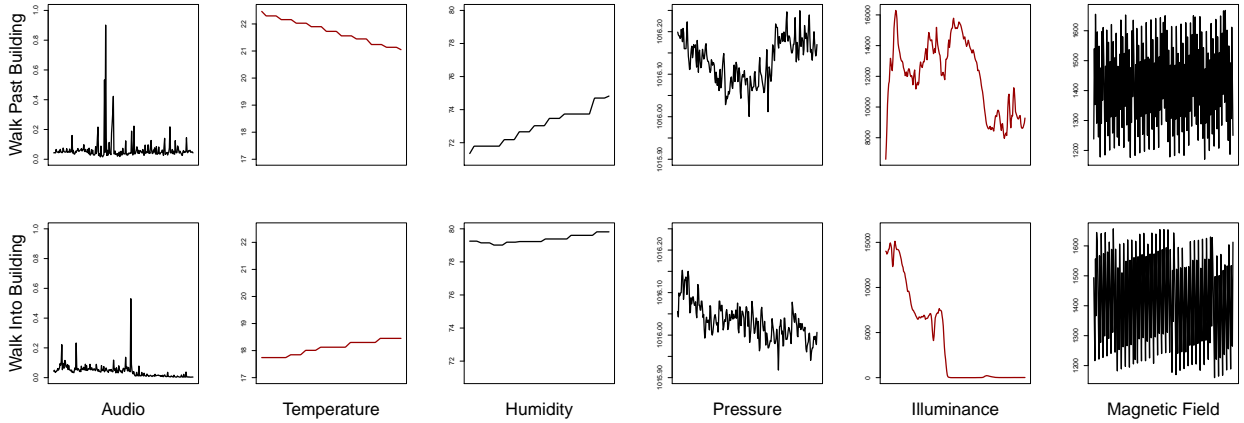
Figure 3: Sensor signatures for two scenarios. The first shows someone walking past a building and the second shows someone walking into the building.

are ranked by the information gain, or the *importance* of each sensor in determining the platial[5] location, i.e., outdoor or indoor, of the device. Illuminance shows the highest value followed by humidity and temperature, indicating that these sensors are the most important for determining a change in platial location from outside of a building to inside. As argued previously, these are important findings as the number of used sensors should be kept to a minimum to reduce battery consumption and ensure that the sensors are not blocked for other applications running on the user's smartphone.

Sensor Type	Gain Value
Illuminance	0.593
Humidity	0.572
Temperature	0.534
Audio	0.391
Pressure	0.260
Magnetic	0.001

Table 1: Information gain of sensors for the indoor vs outdoor experiment

6.1.2 Scenario: Floor Change

A second scenario shows the importance of ambient sensor data in detecting a change in elevation. For indoor spaces (and urban canyons), we know that GNSS cannot provide accurate location information. Detecting and reporting on change in vertical position is not an inherent strength of typical GNSS chips (Yao and Clark, 2000). While Wi-Fi and Bluetooth positioning are offered as solutions to indoor navigation, they are often influenced by positional errors between floors in a building. From an activity-driven perspective, it is important not only to detect a change in the environment, but also to show how a user changes between places (e.g., entering an elevator).

---

[5]We use the term "platial" here in reference to "place-based", similar to how "spatial" refers to "space".
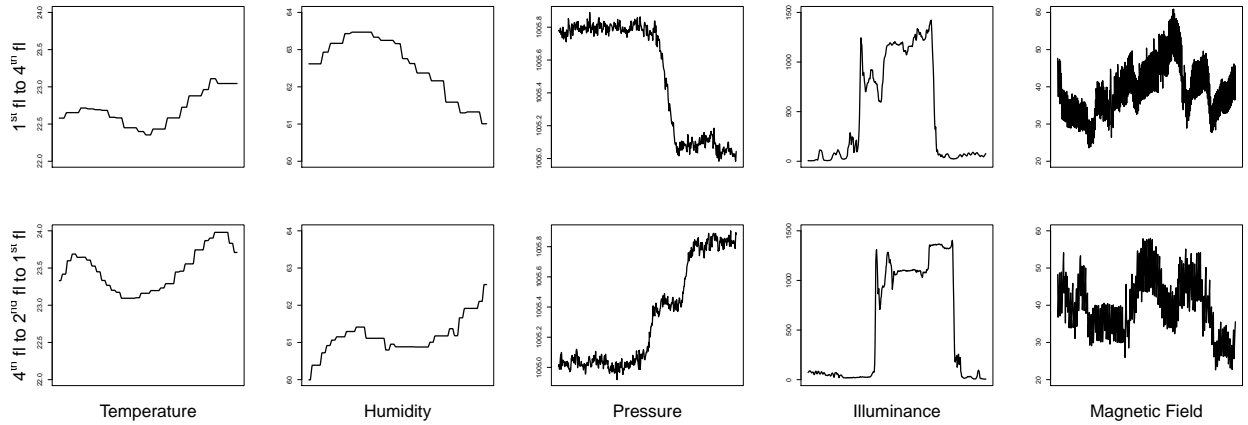
Figure 4: Sensor signatures for entering an elevator on one floor and exiting on a different floor.

Figure 4 shows two scenarios that involve changing floors in a building through the use of an elevator. In the first scenario, a user enters an elevator from the first floor, rides the elevator to the fourth floor and exists on the fourth floor. From a visual inspection of the sensors, we can clearly see separate stages. For example, the illuminance level jumps up by orders of magntiude and the temperature level dips and then recovers over the course of entering and exiting the elevator. The most notable change in sensors is the barometric pressure level dropping with increasing elevation. This analysis is corroborated by evidence in the second scenario in which the user rides an elevator from the fourth floor to the first, stopping briefly at the second floor to allow a passenger to exit. The barometric pressure levels clearly indicate when the second floor stop occurred by the abrupt horizontal section in the graph that stands out from an otherwise steep climb in pressure as the elevator descends.

Sensor Type	Gain Value
Magnetic	0.161
Illuminance	0.674
Humidity	0.505
Pressure	0.303
Temperature	0.345

Table 2: Information gain of each ambient sensor for the binary classification of inside the elevator or outside.

Knowing how much value each sensor contributes to a binary classification can enable us to ignore certain sensors if we are only interested in detecting a particular event. For this scenario, we annotated our data samples by indicating to our device the moment that it crossed the plane of the elevator doors during ingress and egress. This essentially separated samples into two categories, inside the elevator and outside the elevator. The results shown in Table 2 suggest that illuminance and humidity hold the strongest influence for determining whether the device is inside or outside the elevator. The information gain feedback helps us identify elevator ingress and egress events within the building. Being able to detect events is an important feature that helps us improve positioning accuracy by indicating which combination of sensors is most reliable for the given

13

environment, e.g. disabling GPS upon entering a building.

6.2 Place Type Classification

Increasingly, modern devices are beginning to make use of contextual awareness through their automatic responses to certain stimuli such as: entering hands-free mode when a user starts driving their car, bypassing the lock screen at trusted locations, ignoring calls during certain times of the day, and so forth (Baldauf et al., 2007). We recognize an even greater potential to provide both native and third-party apps with contextual information about a user's environment to enable *smart features*. Since much of the measurable context relies on classifying the type of place as well as its activity affordances, we also briefly discuss how to construct a model of place types and how to render the implications.

6.2.1 Mode Switching

One potential use for contextual information about a user's environment is the ability to automatically switch any application's mode to better suit the environment. If a student were to walk into a library and sit down in a study room, their phone could recognize its surrounding environment's activity affordance of *studying* and automatically switch to vibrate mode. Although this may seem trivial to do with geofencing, our approach allows the model to recognize place types and activity affordances in locations that have never before been observed. As long as the ambient sensor signatures agree with the model, geographic location would only be used as a secondary datum to refine the context of a place's environment.

In support of this concept, we ran an experiment to test the accuracy of predicting the activity affordance of a user's environment by using a support vector machine. To train the model, we collected data from ideal study rooms throughout the main library and student center buildings on a University campus at different times of the day. We also collected data from several heterogeneous places around the student center in areas that were deemed not good for studying (e.g, restaurants, social areas, etc.) to ensure the binary classifier would have sufficient samples of alternative observations. In total, we sampled from 12 different places and accumulated more than 740,000 individual observations.

Sensor Type	Gain Value
Humidity	0.577
Illuminance	0.670
Temperature	0.646
Magnetic	0.029
Pressure	0.669
Audio	0.662

Table 3: Information gain of each ambient sensor for the binary classification of studying versus not-studying.

The results of a 10-fold cross-validation yielded an average study affordance prediction accuracy of 99.7%. Each row of training and testing data represent a single observation from all

sensors recorded simultaneously. As seen in Table 3, information gain analysis reveals that illuminance, pressure, and audio levels play the greatest role in determining whether or not a certain environment is suitable for studying. These encouraging results support continued research into the potential uses of ambient sensor signatures. Future work will expand this sample to a wider variety of place types and activities at a larger scale.

6.2.2 Event Annotation

Another practical use for contextual information, such as the type of place, is to annotate events associated with user interactions. Since people are strong at recalling memories by association, a group of device interactions could be annotated with a place and activity type that occurred during its execution. Accurate metadata can enhance the efficiency of searching on documents and events. For example, imagine a traveling scholar reading an article on their tablet while riding a train to the airport. Before they can finish reading the article, they are interrupted by a phone call from a colleague. The scholar later powers down their device, losing the saved state of the PDF viewer they had open. Two weeks later, the scholar queries their device history for what document they had open the moment they received a phone call from their colleague while riding the train (Hu and Janowicz, 2012). In this example, human-computer interactions are grouped into meaningful events so that they can later be processed as individuals, e.g., viewing PDF 'X', receiving phone call from 'Y'. The process of creating events such that they mirror the way a human conceptualizes a group of interactions, along with selecting how to choose and represent relevant contextual information to annotate the events with, e.g., riding the train, is a cognitive research question that is outside the scope of this paper. However, the 'raw materials' of event annotation stem from a knowledge of time and space of the user's immediate surrounding environment. The theoretical framework presented in this paper aims to achieve a similar effect, albeit restricted to the hyper-local spatial domains.

6.3 Activity Detection

A key feature of our platform is to enable user-contributed sensor-driven location-based services. The basis of this idea hinges on the ability to accurately estimate the activity affordances of places. Assuming that the sensor network has a user present at some given place, we can detect the activities that are occurring at that location and then inform a remote user of its current state. To drive the activity detection process, our service collects live ambient sensor readings from a device and feeds them into the service's model to test against previous observations, i.e., place type and activity signatures. To provide more context to our model and improve the accuracy of activity detection, we also include the timestamp of each sample during training and testing phases. Here, we will examine one such scenario using only the ambient sensor data.

Scenario: Finding Places That Suit A Given Activity

A compelling scenario for a student might be to query a smart-campus service by asking it to recommend a place that offers an ideal setting for some activity, such as studying, relaxing, eating, and socializing. To simulate such a service, we ran an experiment in which various places on campus could afford combinations of an established yet subjective set of activity types. Each type

does not preclude any others from occurring (i.e., they are not mutually exclusive). Each activity type either occurs, or it does not. So for example, four activity types yields a total of 16 possible combinations. While our method for collecting the ambient sensor signatures remains the same as in the other experiments presented here, for this scenario we simply show that such data is also fit for geographic services that may query for particular environmental qualities. To ensure that such a service can distinguish between the varieties of sensor signatures that correlate with certain activity types, we selected two activities with high co-occurrence (i.e., eating and socializing) and two activities we expected to have similar, or easily confused, ambient sensor signatures (i.e., studying and relaxing).

We collected ambient sensor signatures with six environmental features (i.e., humidity, pressure, magnetic, temperature, illuminance, and audio) and four activity labels (i.e., eating, relaxing, socializing, and studying) through the GeoTracer App on the mobile devices. After training a multi-class SVM using a 10-fold cross-validation approach (Kohavi, 1995), we generated a classification confusion matrix to evaluate which, if any, activities are likely to be miscategorized and hinder prediction accuracy.

	Observed					
Prediction	eating	relaxing	socializing	studying	other	
eating	3691	0	0	148	0	
relaxing	0	27726	0	27	0	
socializing	2	0	27575	43	6	
studying	0	0	64	14548	64	
Accuracy	0.9987	0.9997	0.9994	0.9921	N/A	**Overall: 0.9961**

Table 4: Confusion matrix for one of the 10-fold cross-validation results in the multi-class SVM classification. The 'other' column accounts for predictions that did not hit their corresponding observations when the original observation had more than one co-occurring tag. These included 'socializing & eating' and 'studying & relaxing'.

As is seen in Table 4, relaxing and studying were not readily confused in our model, although subjectively we felt that the conditions of each environment were similar. Instead we find that studying and eating are most easily confused, indicating that the signatures with those activity labels exhibit similar qualities. Nonetheless, the overall accuracy of our activity-type prediction is quite high (about 99.6%) using the 10-fold cross-validation for evaluating our multi-class SVM classification results. This small-scale evaluation suggests that ambient sensor signatures are suitable data for an activity-type detection service.

6.4 Sensors as Services

An implication to computing activity affordances of place is the ability to approximate the current activity status of a known location (Jordan et al., 1998). For example, a student could tap into their smart-campus sensor network to remotely check the status at some place of interest, e.g. crowdedness of the library. Assuming that a sensor-network contributor is present at the location, their device could provide regular updates to the centralized server, sending ambient sensor samples of the environment to enable up-to-date estimations about the current activity state based on

sensor signatures. This *sensors-as-a-service* approach provides more accurate and real-time information about places as opposed to modern methods that rely on nth-degree information or some preconceived default, e.g., hours of operation.

Scenario: Are the (indoor) basketball courts busy?

For our hypothetical scenario, we chose to ask the question, "Are the basketball courts in-use or not?" Using an SVM with a binary classification of 'in use' or 'not in use', and only sampling when the court was objectively one or the other, we trained our model using samples collected at three different times throughout the day. 235,000 ambient sensor observations were used to train the model, half of them were 'in use' tags while the other half were 'not in use'. We then collected samples to be used exclusively as testing data in order to simulate the challenges that highly dynamic spaces pose against support vector machines. We generated predictions for all 26,000 observations of testing data that were collected from a single trace. The testing data included both 'in use' and 'not in use' tags. The accuracy of our model's predictions against this blind testing data was 85.37%.

To optimize the performance and experiment with improving the accuracy of our predictions, it is important to analyze how much each sensor band contributes to distinguishing between the two different states. To do this, we evaluated the information gain of each sensor.

Sensor Type	Gain Value
Temperature	0.677
Audio	0.675
Humidity	0.670
Illuminance	0.628
Pressure	0.565
Magnetic Field	0.012

Table 5: Information gain for sensor data readings based an active basketball court vs. inactive basketball court.

The results of our information gain analysis, as seen in Table 5, indicate that the magnetic sensor has the weakest correlation to activity status, as one would expect. Interestingly, temperature, audio, and humidity showed the highest information gain values indicating they contributed the most towards determining activity status. These findings are consistent with what one might expect. Temperature and humidity likely increase due to energy exertion and perspiration on the part of the athletes. Given the confines of the closed-door basketball court, it is not surprising that the local temperature and humidity would increase during such an activity. The impact on the audio sensor signature is expected as well as there is considerable difference between an empty court and a basketball game in-progress.

7   Conclusions & Future Work

Modern approaches to developing *smart cities* depend on augmenting existing infrastructure with fixed instruments, an expensive process that requires careful planning in order to provide effective

coverage to its service area. Instruments must be installed adequately enough to deter sabotage yet accessible enough to perform maintenance. Additionally, these node-on-a-network instruments are susceptible to hacking attacks that can lead to compromised networks. In this work, we proposed a mobile sensor-based platform that empowers citizens to collect and share information collected via their mobile devices. This approach allows any environment to become equipped with a free ad-hoc sensor network without depending on proprietary instrumentation. The platform requires only a few hundred milliseconds worth of ambient sensor data to provide several practical service features, sparing users' device batteries while preserving their privacy. And thus the proposed framework could be scaled to a variety of urban areas (e.g., university campuses, shopping malls, city centers, suburbs). Compared to traditional smart city implementations, scaling this framework to larger settings is not limited by physical infrastructure, spatial constraints or financial barriers since it operates by creating an ad-hoc network on top of an existing and dynamic infrastructure. As a proof of concept, we developed a system to collect and share crowdsensing data and showed multiple use-case scenarios about its usages on a university campus.

Future work in this area will involve continued development of the GeoTracer application for data collection, sensor signature training and place-type classification, and scaling both to more place (types) and participants. In addition, further research will explore a temporally weighted approach to sensor value aggregation with the purpose of continually updating place-type classification models with the most recent sensor readings. Lastly, we are currently exploring methods for fully implementing the sensors as a service model allowing ad-hoc querying of hyper-local environmental conditions.

References

Baldauf, M., Dustdar, S., Rosenberg, F., 2007. A survey on context-aware systems. International Journal of Ad Hoc and Ubiquitous Computing 2 (4), 263–277.

Bhattacharya, S., Blunck, H., Kjærgaard, M. B., Nurmi, P., 2015. Robust and energy-efficient trajectory tracking for mobile devices. Mobile Computing, IEEE Transactions on 14 (2), 430–443.

Cardone, G., Foschini, L., Bellavista, P., Corradi, A., Borcea, C., Talasila, M., Curtmola, R., 2013. Fostering participaction in smart cities: a geo-social crowdsensing platform. Communications Magazine, IEEE 51 (6), 112–119.

Carroll, A., Heiser, G., 2010. An analysis of power consumption in a smartphone. In: USENIX annual technical conference. Vol. 14. Boston, MA.

Cerrudo, C., 2015. An emerging us (and world) threat: Cities wide open to cyber attacks. Securing Smart Cities .

Chaudhri, V. K., Cheyer, A., Guilii, R., Jarrold, W., Myers, K., Niekarsz, J., 2006. A case study in engineering a knowledge base for an intelligent personal assistant. In: SemDesk.

Edwards, W. K., Grinter, R. E., 2001. At home with ubiquitous computing: Seven challenges. In: Ubicomp 2001: Ubiquitous Computing. Springer, pp. 256–272.

Elhamshary, M., Youssef, M., 2015. SemSense: Automatic construction of semantic indoor floor-plans. In: Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on. IEEE, pp. 1–11.

Elmaghraby, A. S., Losavio, M. M., 2014. Cyber security challenges in smart cities: Safety, security and privacy. Journal of advanced research 5 (4), 491–497.

Ganti, R. K., Ye, F., Lei, H., 2011. Mobile crowdsensing: current state and future challenges. IEEE Communications Magazine 49 (11), 32–39.

Goodchild, M. F., 2007. Citizens as sensors: the world of volunteered geography. GeoJournal 69 (4), 211–221.

Hu, Y., Janowicz, K., 2012. Improving personal information management by integrating activities in the physical world with the semantic desktop. In: Proceedings of the 20th International Conference on Advances in Geographic Information Systems. ACM, pp. 578–581.

Janowicz, K., 2012. Observation-driven geo-ontology engineering. Transactions in GIS 16 (3), 351–374.

Jordan, T., Raubal, M., Gartrell, B., Egenhofer, M., 1998. An affordance-based model of place in gis. In: 8th Int. Symposium on Spatial Data Handling, SDH. Vol. 98. pp. 98–109.

Keßler, C., Raubal, M., Wosniok, C., 2009. Semantic rules for context-aware geographical information retrieval. In: Smart Sensing and Context. Springer, pp. 77–92.

Klepeis, N. E., Nelson, W. C., Ott, W. R., Robinson, J. P., Tsang, A. M., Switzer, P., Behar, J. V., Hern, S. C., Engelmann, W. H., et al., 2001. The national human activity pattern survey (nhaps): a resource for assessing exposure to environmental pollutants. Journal of exposure analysis and environmental epidemiology 11 (3), 231–252.

Kohavi, R., 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI). Vol. 14(2). pp. 1137–1145.

Komninos, N., 2002. Intelligent cities: innovation, knowledge systems, and digital spaces. Taylor & Francis.

Kuhn, M., 2008. Building predictive models in r using the caret package. Journal of Statistical Software 28 (5).

Lane, N. D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., Campbell, A. T., 2010. A survey of mobile phone sensing. Communications Magazine, IEEE 48 (9), 140–150.

Lin, K., Kansal, A., Lymberopoulos, D., Zhao, F., 2010. Energy-accuracy trade-off for continuous mobile device location. In: Proceedings of the 8th international conference on Mobile systems, applications, and services. ACM, pp. 285–298.

Martins, P. M., McCann, J. A., 2015. The programmable city. Procedia Computer Science 52, 334–341.

Mazilu, S., Blanke, U., Calatroni, A., Tröster, G., 2013. Low-power ambient sensing in smartphones for continuous semantic localization. In: Ambient Intelligence. Springer, pp. 166–181.

McKenzie, G., Janowicz, K., 2015. Where is also about time: A location-distortion model to improve reverse geocoding using behavior-driven temporal semantic signatures. Computers, Environment and Urban Systems 54, 1–13.

McKenzie, G., Janowicz, K., Gao, S., Yang, J.-A., Hu, Y., 2015. Poi pulse: A multi-granular, semantic signatures-based approach for the interactive visualization of big geosocial data. Cartographica: The International Journal for Geographic Information and Geovisualization, The University of Toronto Press 50 (2), 71–85.

Mohan, P., Padmanabhan, V. N., Ramjee, R., 2008. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. In: Proceedings of the 6th ACM conference on Embedded network sensor systems. ACM, pp. 323–336.

OpenSignal, July 2015. Android fragmentation 2015 report. Online, retrieved 02 January 2015.

Paek, J., Kim, J., Govindan, R., 2010. Energy-efficient rate-adaptive gps-based positioning for smartphones. In: Proceedings of the 8th international conference on Mobile systems, applications, and services. ACM, pp. 299–314.

Quinlan, J. R., 1986. Induction of decision trees. Machine learning 1 (1), 81–106.

Santini, S., Ostermaier, B., Adelmann, R., 2009. On the use of sensor nodes and mobile phones for the assessment of noise pollution levels in urban environments. In: Networked Sensing Systems (INSS), 2009 Sixth International Conference on. IEEE, pp. 1–8.

Savelyev, A., Xu, S., Janowicz, K., Mülligann, C., Thatcher, J., Luo, W., 2011. Volunteered geographic services: developing a linked data driven location-based service. In: Proceedings of the 1st ACM SIGSPATIAL International Workshop on Spatial Semantics and Ontologies. ACM, pp. 25–31.

Scheider, S., Janowicz, K., 2014. Place reference systems. Applied Ontology 9 (2), 97–127.

Sen, M., Dutt, A., Agarwal, S., Nath, A., 2013. Issues of privacy and security in the role of software in smart cities. In: Communication Systems and Network Technologies (CSNT), 2013 International Conference on. IEEE, pp. 518–523.

Shannon, C. E., 1951. Prediction and entropy of printed english. Bell system technical journal 30 (1), 50–64.

Suykens, J. A., Vandewalle, J., 1999. Least squares support vector machine classifiers. Neural processing letters 9 (3), 293–300.

Vincent, J., August 2013. London's bins are tracking your smartphone. Independent, retrieved 02 February 2016.

Wang, Y., Lin, J., Annavaram, M., Jacobson, Q. A., Hong, J., Krishnamachari, B., Sadeh, N., 2009a. A framework of energy efficient mobile sensing for automatic user state recognition. In: Proceedings of the 7th international conference on Mobile systems, applications, and services. ACM, pp. 179–192.

Wang, Y., Lin, J., Annavaram, M., Jacobson, Q. A., Hong, J., Krishnamachari, B., Sadeh, N., 2009b. A framework of energy efficient mobile sensing for automatic user state recognition. In: Proceedings of the 7th international conference on Mobile systems, applications, and services. ACM, pp. 179–192.

Xia, C., Schwartz, R., Xie, K., Krebs, A., Langdon, A., Ting, J., Naaman, M., 2014. Citybeat: Real-time social media visualization of hyper-local city data. In: Proceedings of the companion publication of the 23rd international conference on World wide web companion. International World Wide Web Conferences Steering Committee, pp. 167–170.

Yao, H., Clark, R., 2000. Evaluation of sub-meter and 2 to 5 meter accuracy gps receivers to develop digital elevation models. Precision Agriculture 2 (2), 189–200.

Zandbergen, P. A., 2009. Accuracy of iphone locations: A comparison of assisted gps, wifi and cellular positioning. Transactions in GIS 13 (s1), 5–25.

Zheng, Y., Liu, T., Wang, Y., Zhu, Y., Liu, Y., Chang, E., 2014. Diagnosing new york city's noises with ubiquitous data. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing. ACM, pp. 715–725.