

Semantics in Mobile Sensing

Synthesis Lectures on the Semantic Web: Theory and Technology

Editor

James Hendler, *Rensselaer Polytechnic Institute*

Ying Ding, *Indiana University*

Whether you call it the Semantic Web, Linked Data, or Web 3.0, a new generation of Web technologies is offering major advances in the evolution of the World Wide Web. As the first generation of this technology transitions out of the laboratory, new research is exploring how the growing Web of Data will change our world. While topics such as ontology-building and logics remain vital, new areas such as the use of semantics in Web search, the linking and use of open data on the Web, and future applications that will be supported by these technologies are becoming important research areas in their own right. Whether they be scientists, engineers or practitioners, Web users increasingly need to understand not just the new technologies of the Semantic Web, but to understand the principles by which those technologies work, and the best practices for assembling systems that integrate the different languages, resources, and functionalities that will be important in keeping the Web the rapidly expanding, and constantly changing, information space that has changed our lives.

Topics to be included:

- Semantic Web Principles from linked-data to ontology design
- Key Semantic Web technologies and algorithms
- Semantic Search and language technologies
- The Emerging "Web of Data" and its use in industry, government and university applications
- Trust, Social networking and collaboration technologies for the Semantic Web
- The economics of Semantic Web application adoption and use
- Publishing and Science on the Semantic Web
- Semantic Web in health care and life sciences

[Semantics in Mobile Sensing](#)

Zhixian Yan and Dipanjan Chakraborty
2014

[Provenance: An Introduction to PROV](#)

Luc Moreau and Paul Groth
2013

[Resource-Oriented Architecture Patterns for Webs of Data](#)

Brian Sletten
2013

[Aaron Swartz's A Programmable Web: An Unfinished Work](#)

Aaron Swartz
2013

[Incentive-Centric Semantic Web Application Engineering](#)

Elena Simperl, Roberta Cuel, and Martin Stein
2013

[Publishing and Using Cultural Heritage Linked Data on the Semantic Web](#)

Eero Hyvönen
2012

[VIVO: A Semantic Approach to Scholarly Networking and Discovery](#)

Katy Börner, Michael Conlon, Jon Corson-Rikert, and Ying Ding
2012

[Linked Data: Evolving the Web into a Global Data Space](#)

Tom Heath and Christian Bizer
2011

Copyright © 2014 by Morgan & Claypool

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Semantics in Mobile Sensing

Zhixian Yan and Dipanjan Chakraborty

www.morganclaypool.com

ISBN: 9781627053907 paperback

ISBN: 9781627053914 ebook

DOI 10.2200/S00577ED1V01Y201404WBE008

A Publication in the Morgan & Claypool Publishers series

SYNTHESIS LECTURES ON THE SEMANTIC WEB: THEORY AND TECHNOLOGY

Lecture #8

Series Editors: James Hendler, *Rensselaer Polytechnic Institute*

Ying Ding, *Indiana University*

Series ISSN

Print 2160-4711 Electronic 2160-472X

Semantics in Mobile Sensing

Zhixian Yan

École Polytechnique Fédérale de Lausanne (EPFL)

Dipanjan Chakraborty

IBM Research India

*SYNTHESIS LECTURES ON THE SEMANTIC WEB: THEORY AND
TECHNOLOGY #8*



MORGAN & CLAYPOOL PUBLISHERS

ABSTRACT

The dramatic progress of smartphone technologies has ushered in a new era of mobile sensing, where traditional wearable on-body sensors are being rapidly superseded by various embedded sensors in our smartphones. For example, a typical smartphone today, has at the very least a GPS, WiFi, Bluetooth, triaxial accelerometer, and gyroscope. Alongside, new accessories are emerging such as proximity, magnetometer, barometer, temperature, and pressure sensors. Even the default microphone can act as an acoustic sensor to track noise exposure for example. These sensors act as a “lens” to understand the user’s context along different dimensions.

Data can be passively collected from these sensors without interrupting the user. As a result, this new era of mobile sensing has fueled significant interest in understanding what can be extracted from such sensor data both instantaneously as well as considering volumes of time series from these sensors. For example, GPS logs can be used to determine automatically the significant places associated to a user’s life (e.g., home, office, shopping areas). The logs may also reveal travel patterns, and how a user moves from one place to another (e.g., driving or using public transport). These may be used to proactively inform the user about delays, relevant promotions from shops, in his “regular” route. Similarly, accelerometer logs can be used to measure a user’s average walking speed, compute step counts, gait identification, and estimate calories burnt per day. The key objective is to provide better services to end users.

The objective of this book is to inform the reader of the methodologies and techniques for extracting meaningful information (called “semantics”) from sensors on our smartphones. These techniques form the cornerstone of several application areas utilizing smartphone sensor data. We discuss technical challenges and algorithmic solutions for modeling and mining knowledge from smartphone-resident sensor data streams. This book devotes two chapters to dive deep into a set of highly available, commoditized sensors—the positioning sensor (GPS) and motion sensor (accelerometer). Furthermore, this book has a chapter devoted to energy-efficient computation of semantics, as battery life is a major concern on user experience.

KEYWORDS

semantic analytics, smartphone, mobile sensing, crowdsourcing, people-centric sensing, semantic trajectories, trajectory ontologies, semantic activities, activity recognition, energy-efficient computation

Contents

	Acknowledgments	xi
1	Introduction	1
1.1	Mobile Sensing: Definitions and Scope	1
1.1.1	Smartphone-based Sensing	2
1.1.2	Sensing Architectures	5
1.2	Semantics from Sensors	7
1.2.1	Semantic Modeling of Mobile Sensors	8
1.2.2	Semantic Computation from Mobile Sensors	9
1.3	Book Structure	10
2	Semantic Trajectories from Positioning Sensors	11
2.1	The Positioning Sensor	12
2.1.1	Sensor Type and Functionality	12
2.1.2	Dealing with Positioning Data: Technical Summary	14
2.1.3	From GPS to Semantic Trajectories	16
2.2	Semantic Trajectory Modeling	16
2.2.1	Semantic Trajectory Ontology	17
2.2.2	Hybrid Semantic Trajectory Model	22
2.3	Semantic Trajectory Computation	26
2.3.1	Data Preprocessing Layer	27
2.3.2	Trajectory Identification Layer	31
2.3.3	Trajectory Structure Layer	33
2.4	Semantic Trajectory Annotation	37
2.4.1	Annotation with Semantic Regions	40
2.4.2	Annotation with Semantic Lines	41
2.4.3	Annotation with Semantic Points	44
2.5	Summary and Outlook	48
3	Semantic Activities from Motion Sensors	51
3.1	The Motion Sensor	52
3.1.1	Sensor Functionality	52

3.1.2	What can we Learn From Motion?	53
3.1.3	Data Collection	55
3.2	Feature Spaces	57
3.2.1	Data Processing	57
3.2.2	Feature Classes—An Overview	58
3.2.3	Feature Computation and Energy	61
3.3	Activity Learning Techniques	62
3.3.1	Learning Models	63
3.3.2	Choice of Learning Models	67
3.3.3	Testing	67
3.4	Case Study: Micro Activity (MA) Learning	68
3.4.1	Data Collection	68
3.4.2	Feature Vector Representation	69
3.4.3	Results of MA Supervised Learning	70
3.5	Case Study: Complex Activity Learning	71
3.5.1	User Recruitment and Data Collection	72
3.5.2	Data Processing & Sanitization	73
3.5.3	Features for Complex Activities	73
3.5.4	Complex Activity Learning Approaches	78
3.5.5	Learning Performance	80
3.6	Conclusions and Summary	82
4	Energy-Efficient Computation of Semantics from Sensors	85
4.1	Energy-Efficient Mobile Sensing	85
4.1.1	Smartphone Battery Limitations	85
4.1.2	Energy-Efficient Sensing: Hardware Approaches	87
4.1.3	Energy-Efficient Sensing: Software Approaches	87
4.2	Model-based Energy-Efficient Sensing	89
4.2.1	Two-Tier Optimal Sensing	91
4.2.2	Model-Based Optimal Segmentation	93
4.2.3	Model-Based Optimal Sensor Sampling	97
4.3	Energy-Efficient Semantic Activity Learning	99
4.3.1	Classification Accuracy vs. Energy Consumption	100
4.3.2	A3R—A methodology for continuous adaptive sampling	105
4.4	Concluding Remarks	107

5	Conclusion	109
	5.1 Summary	109
	5.2 Challenges and Opportunities	111
	Bibliography	115
	Authors' Biographies	131

Acknowledgments

We would like to express our deepest gratitude to a few excellent researchers and supporters, without whom the book would not exist. Much of the technical work behind this book was conducted when both authors were affiliated with EPFL—Zhixian was first a Ph.D. and then a PostDoc researcher at EPFL, Dipanjan was a visiting researcher. The mobile sensing work was supervised by Prof. Karl Aberer who runs the Distributed Information Systems group at EPFL. We are grateful to Prof. Aberer for his guidance and gracious support to initiate projects in mobile sensing, which formed the fundamental building blocks of this book. This work was sponsored by several Swiss or European research projects at EPFL, including NCCR MICS, OpenSense, OpenIOT, and MODAP.

We owe special thanks to Prof. Stefano Spaccapietra and Prof. Christine Parent at EPFL. They contributed significantly to the key concepts of the semantic trajectories in Chapter 2. Prof. Spaccapietra and Prof. Parent are the first researchers to have initiated the semantic trajectory area, and offered us tremendous support.

Several other researchers and colleagues assisted us in writing this book. Prof. Archan Misra from Singapore Management University researched with us on the topics of activity recognition, which has substantial contributions in Chapter 3. Julien Eberle at EPFL worked with us on collaborative sensing, which becomes a key part of Chapter 4 on energy-efficient mobile sensing. We are thankful to Parikshit Sharma, Sheetal Agarwal, Saket Sathe, and Deepdish Mukherjee who assisted us with reviews and comments.

We would like to acknowledge some of our close family members who have supported us and encouraged us throughout. Dipanjan would like to thank his wife Laura and his parents and in-laws, Joydeb, Nupur, Alexis, and Leela, for their unending love, encouragement, and guidance. Zhixian, in particular, wants to thank his wife Jing and his newborn son Bokai for their patience, support and making the writing of this book a pleasant experience.

We are grateful to the three editors, i.e., Michael B. Morgan (President and CEO of Morgan & Claypool Publishers), Prof. Ying Ding from Indiana University, and Prof. Jim Hendler from Rensselaer Polytechnic Institute for keeping us on track as well as for their active involvement in making this book happen.

Zhixian Yan and Dipanjan Chakraborty
April 2014

CHAPTER 1

Introduction

1.1 MOBILE SENSING: DEFINITIONS AND SCOPE

The term *Mobile Sensing* is not new. Since the late 1980's and early 1990's, one branch of mobile and wireless computing has focused on creating various sensor equipment and establishing sensor networks to monitor phenomena of interest [3, 4] in many applications, such as monitoring atmosphere [45], odor measurement via gas sensor [81], and potholes on road surface [47]. Since then, after the leapfrog of wireless networking technologies, sensors started becoming increasingly connected to each other, or to backend servers via different networking technologies. Usually, each sensor motherboard has a *sensing* module and a *communications* module.

As sensors started becoming untethered, researchers focused on how to increase coverage and reliability of sensing while sensors are moving around. For a long while, traditional wireless network research has focused on different inter-related issues associated with mobile sensors. For example, one rich area of research investigated how to place sensors in an area in order to maximize coverage [25, 39, 78]. Another area looked at how to generate different types of mobility models that can be used by the sensors as “mobility policies” to guide how the sensors should be moving around in an area. A very popular, although old, example is that of random waypoint mobility pattern [22]. The primary application area for this body of work is in battle field surveillance missions, space and land surveillance, deep sea monitoring, etc. Here, the sensor network (i.e., the network of sensor nodes) talk to each other to determine how to best operate, what type of data to collect, when to transmit data to the backend, etc. An overarching need in this body of work is to reduce the energy consumed by the sensors in performing their operations. Energy is a scarce resource, especially in situations where the node is dependent on a battery for its survival.

While the wireless sensor network community focused on the above application areas and associated technical problems, another body of research work started surfacing up around the mid 1990's. This was called *wearable computing*. In this book, the primary objective was to monitor and map movement patterns of living organisms—animals, birds, humans, etc.—using a plethora of sensors that are worn or carried with the organism. More often than not, the primary direction of research in wearable computing has been *inward*, i.e., to use sensors to make observations about the carrier, rather than to sense a phenomenon of interest (e.g., environment). The observations made have been used in innovative ways to make inferences about the carrier [108, 128] and to establish argument reality [149]. For example, on-body acceleration sensors, worn in different parts, can be used to figure out locomotory movements, location sensors can be used to figure out

2 1. INTRODUCTION

mobility patterns, and RFID badges can be used to figure out proximity to certain surroundings, e.g., Bluetooth enabled “smart objects” in a smart room.

The line between *mobile (moving object) sensor networks* and *wearable computing* started getting blurred around the mid 2000's with the advent of smartphones. Smartphones come with a number of special-purpose sensors (e.g., GPS, Accelerometer, Gyroscope) or communication units (e.g., Bluetooth, WiFi). The output trace from these sensors or communication units reveal rich information about the smartphone, as a new paradigm of mobile phone sensing [96], which is also called *people-centric sensing* [23]. A smartphone can be easily seen as a powerful computing and communication device that can be used with traditional sensors (e.g., temperature sensor) to report statistics about underlying environment. For the traditional sensor networking community, a smartphone presents a powerful, pervasive, and well-adopted computing and communication tool that can be used in many different ways. For the wearable computing community, the smartphone represents a powerful wearable multi-sensor unit that can be used to infer knowledge about the carrier as well as the underlying surrounding. With the rapid growth of cellular data and network standards, we are witnessing an improvement in the data bandwidth available with us. Scientists are of the belief that this combined trend of increased network bandwidth and speed, coupled with the advanced sensing and processing capabilities of smartphones and its ability to operate as a *sensor docking station*, is going to change the landscape of mobile sensing in the future.

This book goes beyond the contemporary meaning of the term mobile sensing and investigates semantic data extraction in an era where traditional sensor networks have united with wearable sensing. In particular, we are largely going to focus on the area encompassing (1) embedded or extended sensors that are on the smartphone, (2) carriers that are usually people or community-driven objects such as cars, buses, etc., and (3) mobility that is usually un-organized and un-structured and bottom-up, i.e., the community autonomously governs the mobility. Hence, our scope of the term *mobile sensing* is going to largely represent the above real-life conditions.

In the next subsection, we provide a high-level background of smartphone sensing, sensing architectures, and trends.

1.1.1 SMARTPHONE-BASED SENSING

Andrew Campbell from Dartmouth College, Hanover, New Hampshire, USA, is one of the first researchers to have coined the term “people sensing” back in 2005 [24]. The word meant that traditional sensor networks are going to slowly evolve into a model where people will be carrying sensing devices and using personal devices to sense heterogeneous phenomenon of interest—in particular, on the applications of urban sensing. However, in 2005, smartphones were still difficult to program. The capabilities in terms of processing power and on-board sensing units were also limited. Since then, over the last eight years, significant progress has been made in the market. The introduction of the iPhone OS¹ in 2007 and the Android² platform in 2008 opened up the

¹<http://www.apple.com/iphone/ios/>

²<http://www.android.com/>

operating system for rapid and easy programming and community-driven development. Progressively with time, more and more cheap and easy-to-program sensors started getting embedded into the phone. The phone that we carry, slowly but steadily, started moving from a perception of *communication device* to a *personal intelligent assistant*.

The smartphone today has several on-board sensing units. Many physical sensors are those units that explicitly measure certain attribute of the environment (e.g., location, temperature). Other sensors can be those whose primary functionality is not to sense, but the data implicitly contains some traces of the environment. An example of an “explicit” physical sensor is the Global Positioning System (GPS) sensor. An example of an “implicit” sensor is Bluetooth. Bluetooth is primarily used for peer-to-peer communications, but the device discovery module of Bluetooth can periodically report the other Bluetooth-enabled devices that can be observed in its vicinity. This can be used to infer some context about the surroundings of the device. We provide a typical list of physical sensors that come embedded with some of the latest smartphone handsets today.

- **Global Positioning System—GPS.** This is one of the most well-known sensors on the smartphone. It computes the position of the smartphone using satellites and their ground stations as a reference frame. The output is typically a $\langle \textit{latitude}, \textit{longitude}, \textit{altitude} \rangle$ tuple, along with other information such as velocity and direction. An extensive tutorial on GPS is provided by Trimble,³ and more detailed explanations of GPS can be found in [125].
- **Accelerometer—ACC.** An accelerometer records acceleration along three mutually perpendicular axes x, y, z . When the smartphone is placed horizontal on a table, one axis points towards gravity (\vec{g}) vector. Accelerometers play an active role in measuring physical activities [156]. They are typically used to detect the orientation of a smartphone and widely used today in various gaming and activity recognition applications [139, 170]. Readers interested in knowing physical properties of a typical accelerometer chip may refer to wikipedia.⁴
- **Gyroscope—GYR.** A gyroscope is used to measure orientation of the smartphone using principles of angular momentum. It can measure the rate of change of angle along a particular axis and is typically reliable to detect short-term and bursty orientation changes accurately. Smartphones typically use a combination of accelerometer and gyroscope to calibrate each other [37] as well as reliably detect orientation changes, locomotive states, etc. [107]
- **Bluetooth—BT.** Bluetooth is used to exchange data over short distances between two peers using radio transmissions in the ISM band of 2400–2480 Mhz. Bluetooth can hence discover other peer Bluetooth-enabled devices around itself. Thanks to the new Bluetooth LE (Low Energy) technology, smartphones can more efficiently capture proximity data. This property has been widely used in research to estimate and sense real-time social surroundings, e.g., estimating the population density [161], computing the groups [41], and inferring social contexts [142].

³http://www.trimble.com/gps_tutorial/howgps.aspx

⁴<http://en.wikipedia.org/wiki/Accelerometer>

4 1. INTRODUCTION

- **Light—LT.** The light sensors are typically photodiodes measuring light intensity, by correlating it with measured current. Light sensors are widely used to control screen brightness in different luminous settings. It is also used to lock the touchscreen when a user is holding the phone against his ear. Light sensor is also integrated with other sensors for intelligent services such as better localization with the combination of sound [7] and phone position detection (e.g., in pocket, on table, in hand) together with accelerometer [144].

The list of phone sensors does not end here. There are other sensors such as NFC-related proximity sensors and barometers. Of course, the camera and the microphone—two most prevalent components can also act as sources to feed information about a smartphone’s surrounding, visual, and acoustic context. Apart from these embedded sensors, smartphones come with input/output interfaces (like micro-USB ports) that can be used by many accessory manufacturers to plug-in accessories like a headset, etc. Recently, a startup built accessory that plugs into the mobile devices’ (iPhone or Android devices) audio jack for flexible credit card payment using Square.⁵ Researchers and practitioners are investing money to prepare novel accessories which can be used as add-on gadgets along with a smartphone.

Each of the above sensors have specific utility for their inclusion in the handset or as an accessory. For example, GPS is used to supply real-time location of the device to map applications (such as Google maps), which can help the user to localize herself in unknown surroundings, find directions, etc. An accelerometer measures the acceleration of the device, typically along three dimensions—along the vector pointing towards the center of earth (called the gravity vector), and two vectors perpendicular to the gravity vector. Accelerometer is often used as a motion sensor, for stabilization of the camera during videography and photography. It is also used to measure the orientation of the phone. Similarly, proximity and light sensor is often used to control the illumination intensity of the screen and save on power.

While the first wave of applications of these sensors is primarily to drive user experience (e.g., to change display orientation by analyzing real-time accelerometer data), researchers across the globe are focusing on a subsequent wave of innovative applications, exploiting these smartphone sensors to capture, analyze, and predict several environmental properties of the users. Lane et al., in their survey on *Mobile Phone Sensing* in 2010 [96], observed “Now phones can be programmed to support new disruptive sensing applications such as sharing the user’s real-time activity with friends on social networks such as Facebook, keeping track of a person’s carbon footprint, or monitoring a user’s well being. Second, smartphones are open and programmable. In addition to sensing, phones come with computing and communication resources that offer a low barrier of entry for third-party programmers (e.g., undergraduates with little phone programming experience are developing and shipping applications). Third, importantly, each phone vendor now offers an app store allowing developers to deliver new applications to large populations of users across the globe, which is transforming the deployment of new applications, and allowing the collection and analysis of data far beyond the scale of what was previously possible. Fourth, the mobile

⁵<https://squareup.com/>

computing cloud enables developers to offload mobile services to back-end servers, providing unprecedented scale and additional resources for computing on collections of large-scale sensor data and supporting advanced features such as persuasive user feedback based on the analysis of big sensor data.”

In [154], the authors categorize people-driven sensing using physical sensors in terms of three types of observation properties: (1) spatio-temporal properties (e.g., presence, count, location, track, identity); (2) behavioral properties (e.g., pose, action, activity, behavior, group behavior); and (3) physiological properties (e.g., temperature, blood pressure, heart rate). In addition to sensing these personalized properties, sensors also capture information of the nature, e.g., to monitor environment like air quality and climate change, as well as their consequences [29]. Many novel sensing applications have been emerging, such as using mobile phones to detect earthquakes [51], deploying sensors on public transports for air quality monitoring [2], and the Copenhagen Wheel for promptly analyzing pollution levels, traffic congestion, and road conditions using e-bikes.⁶

It is worth noting that there is more emerging interest in studying soft-sensors rather than previously mentioned physical sensors. For example, human can easily report massive real-time social data using the rapidly growing social networks, e.g., Facebook, Twitter, YouTube, and location-based social networking services like Foursquare check-in. Therefore, in the “human as sensors” paradigm, a sensor is not necessarily a hardware sensor but also a virtual sensor also-called a “social sensor” or “logical/soft sensor” in literature. Such virtual sensing in social networks allows efficient and effective information sharing and propagation, with an unprecedented deployment scale, which in turn open novel data-driven applications in numerous domains such as health, transportation, energy, disaster recovery, and even warfare [147]. These applications are often referred to as “crowd-sourcing” or “citizen sensing” scenarios, since they are capitalizing on the power of crowds and relying on a large scale of citizens participation [17]. In addition to using public social networks/media like Twitter as virtual sensors in building crowdsourcing applications, several dedicated open-source platforms are alternatively built for easy crowdsourcing development, e.g., GeoCha⁷ and Ushahidi.⁸ It is worth noting that user reports via virtual sensors would be biased and subjective, but we can use that fact to identify communities or group and do relevant data filtering. In this book, the main focus is on studying the semantic computation over physical sensors, not soft sensors.

1.1.2 SENSING ARCHITECTURES

Information sensed by the users and their smartphones may be transmitted to a back-end server. Information from multiple devices and users may be combined together to reveal significant trends of an environment like predicting air quality, crowd levels, and generally, information

⁶<http://senseable.mit.edu/copenhagenwheel/>

⁷<http://instedd.org/technologies/geochat/>

⁸<http://ushahidi.com/>

6 1. INTRODUCTION

relating several important city management sectors like traffic issues, neighborhood issues, emergency, civic complaints, etc. These developments are known under the names of “urban sensing” [36, 95], “participatory sensing” [21, 63], “opportunistic sensing” [85], “community sensing” [90], “crowdsensing” [60], “crowdsourcing” [79, 89], “people sensing” [24], etc. These buzz words all describe the space of sensing architectures from various application perspectives. Nevertheless, the key point from all buzz words is to build applications that can benefit from participators or volunteers of a large number of people with mobile devices. The data/contributions from participators can be generated not only actively but also passively. The fundamental difference between this new paradigm of community-driven sensing and traditional sensor networks is in its organic involvement model, and autonomous, human driven nature of the sensor network.

Among these buzz words, *crowdsensing* and *crowdsourcing* are largely used, in recent years, for describing such community-driven mobile sensing designs. Crowdsensing is a relatively new term compared to Crowdsourcing. The fundamental difference between crowdsourcing and crowdsensing is in the nature of the task. Typically, crowdsourcing tasks are top-down driven, i.e., a task is decomposed into suitable pieces that can be executed by human endpoints. Thereafter, different solutions are used to reach out to the community and get the pieces executed. Crowdsensing, on the other hand, is a bottom-up driven method, i.e., the community is organically involved in sensing a phenomenon of interest. The basic blueprint of a crowdsensing application consists of a sensing agent running on the phone, while a back-end server aggregates results and provides services.

In terms of on-the-field deployments, there are quite a few trends. With the growing number of cities and populations within them, continuous monitoring of city’s infrastructure for growth and sustainable development is gaining more and more importance. Increasingly, the city is being instrumented and inter-wired with millions of heterogeneous sensing infrastructures (traffic surveillance, CCTV footages in popular areas, sensors in malls, indoor localization). In parallel, the community is increasingly starting to play a role in monitoring urban dynamics using their smartphones, powered by sensors and social media tools. Therefore, the scope of “human as sensors” extend beyond the embedded sensors and also include what can be sensed by a user’s biological sense organs. Social media offers a “channel” for capturing such sensed data about a city, its traffic, social and public events, etc. Social media data collection infrastructures, combined with a handy smartphone with several social update applications, provide an easy-to-use mechanism for users to provide just-in-time updates, and capture unprecedented data about a city. For example, we already know how social media assisted in the Spring uprisings in Africa. Each of these social media infrastructures largely follow a client-server oriented architecture, where a thin client is used to provide updates to a user and receive inputs. There is a significant amount of work happening in the social media space to glean relevant knowledge from this high volume, noisy, text data stream [74, 84].

The concept of “humans as sensors” to observe, infer, and predict properties of our environment has gained significant attention over the last few years. For example, GPS trajectories

generated by running the GPS sensor periodically from smartphones can be used to understand urban densities and population dynamics. At a personal level, such trajectories contain rich information of our day-to-day life and can be used in several ways to provide us with meaningful recommendations, exploit activities and activity logs, and better understand health quality. Similarly, accelerometer can be used to monitor the amount of time we are walking or performing fitness routines. Research has made significant progress on understanding algorithms—methodologies for knowledge extraction from these sensor streams. However, quality and reliability of predictions “in the wild,” i.e., as the user moves around in day-to-day life is an open question. It has been studied that laboratory-constrained environments and results obtained via such experiments are not synergistic with experiments in the wild. In this book, we will provide the reader with a flavor of algorithms and techniques that are used for knowledge extraction from such smartphone sensor data streams, and provide the reader with inputs and insights on open issues and challenges of applying these in real-world settings. The learning and inference should be robust and domain specific.

In this book, we would discuss and present the emergent architectures that are being proposed to capture data and extract knowledge or semantics out of such city-scale deployments. In particular, we focus on the main techniques for inferring meaningful information from sensor data captured through mobile phones towards such a model of crowdsensing. In addition, we present energy-efficient sensing approaches that are highly applicable in smartphone platforms.

1.2 SEMANTICS FROM SENSORS

The word *semantics* popularly means “the study of meanings.” In linguistics, it is used to represent the interpretation of words, phrases, signs and symbols in a given context (i.e., a bag of such words, phrases, symbols) [150]. Compared with its extensive study in the area of Web and natural language processing, in particular with the booming of Semantic Web technologies in recent years, *semantics* is much less studied in the mobile sensing area. For mobile sensing, it’s non-trivial to provide a universal definition for semantics; but, in general, *semantics* in mobile sensing very broadly refers to any meaningful knowledge that we can extract from the sensor data. For example, GPS sensor provides a geometric location point with the $\langle \textit{latitude}, \textit{longitude} \rangle$ coordinates. The semantics of this point could be more meaningful geographical information. Sources like maps could be used to extract where the GPS point is lying, e.g., Is it on a road? Is it in a residential area? Is it in a lake?

In this book, the objective is to extract the semantics (meaningful information) from raw mobile sensing data generated by smartphones, for better understanding the phenomena in the mobile sensing era. In the subsequent chapters, we will select some representative sensors (e.g., GPS, accelerometer) in the mobile phone, and present techniques for identifying semantics of these sensors. For each sensor type, we will design semantic modeling of the sensor data, which can provide different levels of abstraction of sensor semantics; hereinafter we accordingly present

and summarize the key supporting techniques to compute the high-level semantics from the low-level raw sensor data stream.

1.2.1 SEMANTIC MODELING OF MOBILE SENSORS

In Section 1.1.1, we summarized a rich list of physical and virtual sensors in the people-centric sensing domain using today's smartphones. These sensors include GPS, WiFi, GSM, accelerometer, gyroscope, barometer, Bluetooth, light, acoustic, etc. In this book, we focus on semantically analyzing three representative sensors that play a significant role in mobile sensing, i.e., *positioning sensors* (like GPS, WiFi) and *motion sensors* (like accelerometer, gyroscope). The fundamental step for extracting semantics from these sensor data is to provide a semantic model to represent sensor streams. In this book, two important modeling techniques will be extensively discussed for modeling sensor data, i.e., “conceptual modeling” and “ontologies.”

- **Conceptual Modeling.** Conceptual modeling has been extensively used for formally describing some aspects of the physical and social world around us for the purpose of understanding and communications [16, 116]. It has already played a success role in building conventional information systems. Conceptual models have also been actively used in representing sensor data, e.g., the “stop,” “move,” and “trajectory” concepts for modeling the abstraction of GPS coordinations [124, 146]. In [152], the authors summarize a set of conceptual models that can be used in sensors in the ubiquitous computing paradigm, particularly focusing on context awareness-oriented applications.
- **Ontologies.** In the Semantic Web, ontologies is the building-block technology for modeling the concepts within a domain, providing “*formal, explicit specification of a shared conceptualization*” [12, 65], etc. Ontologies also plays an important role in modeling sensors and their high-level semantics, e.g., trajectories for GPS data [166], activities for accelerometer data [141], context modeling and reasoning [159], etc.

In line with these conceptual models and ontologies, we present a summary of key semantic modeling approaches for representing sensor data in mobile sensing, and provide formal definitions that could be used for extracting meaningful semantics from raw sensor data. In this book, as we focus on two representative sensor types, the dedicated semantic models will be the following two aspects.

- **Semantic Trajectories.** We can study a sequence of such GPS points of a moving object to infer the set of geographical attributes the moving object is passing on its way. The raw GPS sequence (trajectory) can be represented as a sequence of meaningful *stops* and *moves*, which can be deduced by the velocity of the moving object. Furthermore, the trajectory can be enriched with geographic attributes, e.g., in the morning, the user was moving from home to office, and during the noon time, the user moved from office to a nearby restaurant for lunch. All these are examples of semantics from GPS data, so-called “semantic trajectories” [164, 167].

- **Semantic Activities.** Let us take accelerometer as another example. The raw data of the accelerometer primarily provides the three-axis acceleration components of the device. However, using appropriate machine learning algorithms, scientists have discovered that the readings can be used to reliably infer locomotive states of a user. Examples of such locomotive states are: walk, sit, stand, jump, run, etc. In wearable sensing literature [66, 88], accelerometer worn in multiple body parts can be used to recover fine-grained limb movements. If we take a sequence of such readings over time, a combination of such limb movements can provide us with vital information to infer *complex activities* like making coffee, running on a treadmill, having lunch, etc. These complex activities are often referred to as *activity routines* [13] or high-level *semantic activities* [137, 163] in contrast to the aforementioned locomotive states (also called “micro-activities”) such as sit, stand, walk, jog, and drive.

1.2.2 SEMANTIC COMPUTATION FROM MOBILE SENSORS

On one side, this book presents semantic modeling of various sensors in smartphone-based new era of mobile sensing, and particularly focuses on two types of sensors and establishes the dedicated semantic concepts, i.e., “semantic trajectories” and “semantic activities” that were briefly mentioned in the previous section. On the other side, the book provides the reader with an overview of algorithms that are used to progressively enrich a set of sensor data points, coming from smartphone and other embedded sensors, to retrieve these defined semantics corresponding to the data.

In contrast to top-down conceptual and semantic modeling of sensor data representation in the mobile sensing era, we additionally provide extensive summaries of key techniques to compute these semantics from the raw sensor data from a bottom-up viewpoint. Dedicated algorithms about sensor data acquisition, preprocessing (e.g., data cleaning, compression), and statistical modeling will be discussed in details in the book. Many related data mining and machine learning algorithms can be applied into these semantic computations of sensor data. Moreover, high-level semantic annotations, query processing, and visualization of computed semantics are also important facets towards the new era of mobile sensing.

It is worth noting that energy consumption (battery drain) is one of the key issue in smartphone-based mobile sensing. Therefore, we will also provide a dedicated chapter to discuss energy-efficient semantic computation techniques.

We would like to bring up a very important question: Do semantics extracted from mobile sensing reveal private data of users? For example, it has been well studied [20] that the top two significant locations of most users having a city-centric lifestyle reveal their home and their office areas. It is very important that applications and infrastructures designed to “consume” such semantic data preserve privacy of users and provide appropriate tools for users to protect how much personal information flows out and to which stakeholders. Real-world adoption of mobile sensing applications would need to have effective algorithms for extracting semantics as well as first-class

privacy preservation algorithms. Thankfully, there has been a significant amount of work in the area of privacy preservation of mobile data. Privacy-preserving data mining [58, 99] investigates how to design new data mining algorithms, while ensuring privacy of users. This book would not cover the topic of privacy, but focus on the knowledge and semantics extraction algorithms from mobile sensing data.

1.3 BOOK STRUCTURE

In the forthcoming chapters, we focus on presenting techniques for knowledge extraction from smartphone sensors, computing semantics of citizen-sensed data, and representation of the sensor semantics at various levels.

In Chapter 2, we present methods of semantic enrichment of location streams captured from GPS sensor of the smartphone, and build so-called semantic trajectories. The key challenge here is to address the inherent unreliability of data sources (e.g., GPS weak signal and noisy data) and to extract various semantic informations (e.g., frequent visit places or mobility behaviors).

In Chapter 3, we discuss methods of recognizing semantic locomotion and activity patterns (named “semantic activities”) from acceleration streams, captured from accelerometers of smartphones. Several simple locomotive states (like walk, sit, stand) and complex activities of daily living (including office work, lunch, home cooking, and relax) can be extracted from such motion streams.

In Chapter 4, we address an important challenge of computing semantics from mobile sensors, i.e., the limitation of smartphone battery capacity. A set of state-of-the-art technologies for energy-efficient semantic computation over mobile phone sensors will be summarized. In particular, we will focus on energy-efficient sensing and activity recognition using smartphones.

Finally, Chapter 5 concludes the book and points to other important research and technical challenges in computing semantics in mobile sensing, such as privacy-preserving semantic extraction, and the combination with the emerging platforms of cloud computing as well as mobile social networks.