Trust and Privacy in Mobile Experience Sharing - Future Challenges and Avenues for Research

Ioannis Krontiris*, Marc Langheinrich† and Katie Shilton‡
*Institute of Business Informatics, Goethe University Frankfurt, Germany; ioannis.krontiris@m-chair.de
†Faculty of Informatics, University of Lugano (USI), Lugano, Switzerland; marc.langheinrich@usi.ch
‡University of Maryland, College Park, MD, USA; kshilton@umd.edu

***Abstract*—Mobile consumer devices are increasingly used as personal sensing instruments, where users record their daily habits, track their physical activity, or monitor their health. Research is underway to extend today's diversity of vendor-designed "walled garden" repositories, ultimately repositioning individuals as producers, consumers, and remixers of a vast openly shared public data set. By empowering people to easily measure, report, and compare their own personal environment, such tools transform everyday citizens into reporting agents who uncover and visualize unseen elements of their own everyday experiences. With this important new shift in mobile device usage – from a communication tool to a ubiquitous "experience sharing instrument" – comes a new dimension in trust and privacy challenges. Today's privacy and trust tools that address Web surfing and simple location-based services already struggle to be adopted in practice. We argue that in order to prepare for tomorrow's sensor sharing, privacy and trust must be addressed holistically, incorporating both technical approaches and actual sharing behavior. This article summarizes the results of a five-day Dagstuhl Seminar on mobile experience sharing and outlines future research necessary in this domain.**

I. Introduction

The increasing availability of sensors on today's smartphones and other everyday devices, carried around by millions of people, has opened up diverse kinds of information gathering by people and their devices. Eventually, researchers envision the creation of a unified data-sharing infrastructure, where people and their mobile devices provide their collected data streams in accessible ways to third parties interested in integrating and remixing the data for a specific purpose. This trend is often named *Mobile Crowdsensing* [1].

Mobile crowdsensing applications can span a broad spectrum of subjects, such as public health and wellness, environmental monitoring, sustainability, urban planning, and cultural expression. A popular example is using sensor-equipped mobile phones that allow everyday citizens to collect environmental data such as noise and air pollution, in an effort to improve both their individual lives (e.g., by mapping pollution-free bike routes) and their community (e.g., by alerting city officials to excessive traffic levels in a neighborhood). Moreover, by tracking one's own activities – such as work, sleep, exercise, diet, or mood – and sharing this information with other people, "quantified self" scenarios further offer novel applications and opportunities for self-improvement [1], [2].

It is obvious that these novel paradigms open up a wealth of concerns for personal privacy: gathering and sharing one's own activities allows others intimate insights into one's work performance, health status, physiological development, professional and leisure activities, and even psychological well-being. Even innocuous data such as temperature readings and carbon monoxide levels may give away personal information – in the form of location coordinates and timestamps that indicate an individual's movement over days and weeks – and can easily lead to unwanted repercussions, e.g., by lowering real estate value in neighborhoods that record unfavorable sensor readings. Moreover, as sharing practices become more fluid than in desktop-based online environments, control over such information flows becomes even harder to maintain [2].

During the last several years, research on Privacy-Enhancing Technologies (PETs) has produced a wide variety of mechanisms [3]. Yet for all innovation in this field, PETs have so far not been widely adopted in practice. It is hard to provide a definitive explanation for this, given the complexity and interdisciplinary nature of the problem. However, various scholars have repeatedly identified the lack of adequate consideration of the various stakeholders' viewpoint as one of the core issues and they point out that simple technological solutions are often not effective in capturing the interests and concerns of users [4], [5]. Consequently, research in mobile crowdsensing – i.e., the development of tools and architectures to enable the ubiquitous sharing of experiences – requires a renewed effort in providing corresponding tools for ensuring trust and privacy in such infrastructures and applications.

This article summarizes the output of a recent Dagstuhl Seminar on this topic, titled "My Life, Shared – Trust and Privacy in the Age of Ubiquitous Experience Sharing" [5]. The seminar brought together experts in computer science and engineering, economics, social sciences, and legal sciences, in order to discuss latest developments and future challenges for trust and privacy in mobile crowdsensing. In particular, the participants identified challenges on three distinct levels, which we will highlight in this article:

- Privacy Engineering: We highlight research challenges of integrating PETs in mobile crowd sensing applications.
- Sharing Practices: Taking the user's perspective, we investigate how the disclosure preferences and practices of sharing personal sensitive data can change and how our tools can support this process.
- Fairness and Social Justice: We might tend to think that the development of new application in the area of

mobile crowd sensing will have only positive impact on the society, but there might be also negative ones. In particular, we discuss the broader concerns about the impact of sensing on social justice, and how this might be impacted – positively or negatively – in a variety of social domains.

Further details can also be found in the seminar report cited above – here we make an effort to cover the topic broader and present a high-level overview based on the three identified levels. We explicitly acknowledge the contributions of the individual seminar participants to each of these areas at the end of this article.

II. Privacy Engineering

A. Architecture Approaches

Several of the works on mobile crowd sensing systems started differentiating very early between two data collection models. In the participatory model, users are actively involved in the collection process by deciding on the spot when to report data, while in the opportunistic model, sensor sampling occurs whenever the state of the device (e.g. geographic location) matches the application's requirements described in a sensing task, without the knowledge of the individual phone user.

Independently from the collection model however, what is common in the majority of existing architectures is that the sensing data collected from the mobile phones are stored in centralized servers, creating massive databases of individuals' location, movements, images, and even health data. This is depicted in Figure 1(a). After collecting the data, the entity controlling the database aggregates, processes and represents them through various interfaces (e.g. statistical data on a map).

This paradigm raises several challenges concerning information access and reciprocity. Who controls data collection and who owns the data or benefits from them? In cases where users collect data to be used by service providers, this data is collected, stored, and analyzed by those providers typically out of view of the individual whose life they describe. The collection of the data is not always restricted to the purpose for which they were collected. Also, individuals cannot pose restrictions on the collection and processing of their data and the data collected from them are not made available back to them through proper interfaces.

To deal with the power imbalance created in such paradigms, architectures taking a more user-centric approach started to appear. What these architectures try to do is to enable individuals to supervise and limit personal data disclosure and exercise access control to their data by third parties. Several existing solutions in crowd sensing applications (see for example [6], [7]) suggest a vault-like entity to provide an online trusted storage and processing. Mobile phones sense and upload data to this vault, which is owned and controlled by the individual. This is shown in Figure 1(b). The process of storing personal data streams is decoupled from the sharing of that information. After the collection and archival of data, the users can define their own privacy policies and review/control who can see which kind of data.

A similar concept is the Personal Data Service (PDS), which was one of the topics discussed amongst the participants in our Dagstuhl Seminar [5]. The PDS is a trusted container for aggregating, storing, processing and exporting personal data, but it is extended to include all data regarding the user, either they are user-generated or they are obtained from other sources, e.g. service providers, including personal data collected and published by third parties. Users are in control of all data stored in the PDS, which includes the option to share or sell parts of this data. In addition to storing data, the PDS can execute code to process this data locally, filter them and apply privacy-preserving methods like obfuscation or generalization of sensitive personal information.

One of the benefits of this approach is that it increases transparency, awareness and engagement of users with their data and gives them an opportunity to validate the integrity of their data. It also supports domain specific identity management, so that individuals can configure appropriate roles and associated data to be shared with thirds party services.

However, there are still several issues to be investigated and real challenges that need to be addressed. For example, what kind of incentives would be needed for the initial data providers to engage and open up the personal data APIs that are needed to fuel the PDS and associated applications?

There are also several open questions with respect to privacy. Even though the PDS can increase transparency, awareness and engagement of users with their data, it is neither obvious nor guaranteed that the PDS will resolve user privacy problems and several of them remain open. The PDS, for instance, may provide users with local control over their data, but that might not prevent third parties from collecting and exploiting user data. It might even facilitate them, by functioning as a central repository for the user's diverse data streams, depending on how locally derived data will be exported to third-party application providers [5].

B. Integration of PETs

1) Large-Scale Integration: Over the last 20 years, the privacy community has developed a large pool of tools and primitives for solving various aspects of the privacy problem. Yet they are not widely adopted in practical systems today. There are several reasons for this, since privacy is a complex and interdisciplinary issue. An earlier Dagstuhl Perspectives Workshop reported on several dimensions of this problem and gave suggestions for how to address the challenges ahead [4]. One of the problems it identifies is that the integration of the individual technological tools in large scale systems and their interaction with each other is not well understood.

Indeed, there is little experience on how privacy-enhancing technologies scale, when they are deployed on large and open-ended networks with decentralized control and governance structures, large populations, and qualitatively different scales of data collection and processing. As it is especially evident from the crowd sensing paradigm, data collection practices and data flows evolve rapidly, and empirical data about this evo-

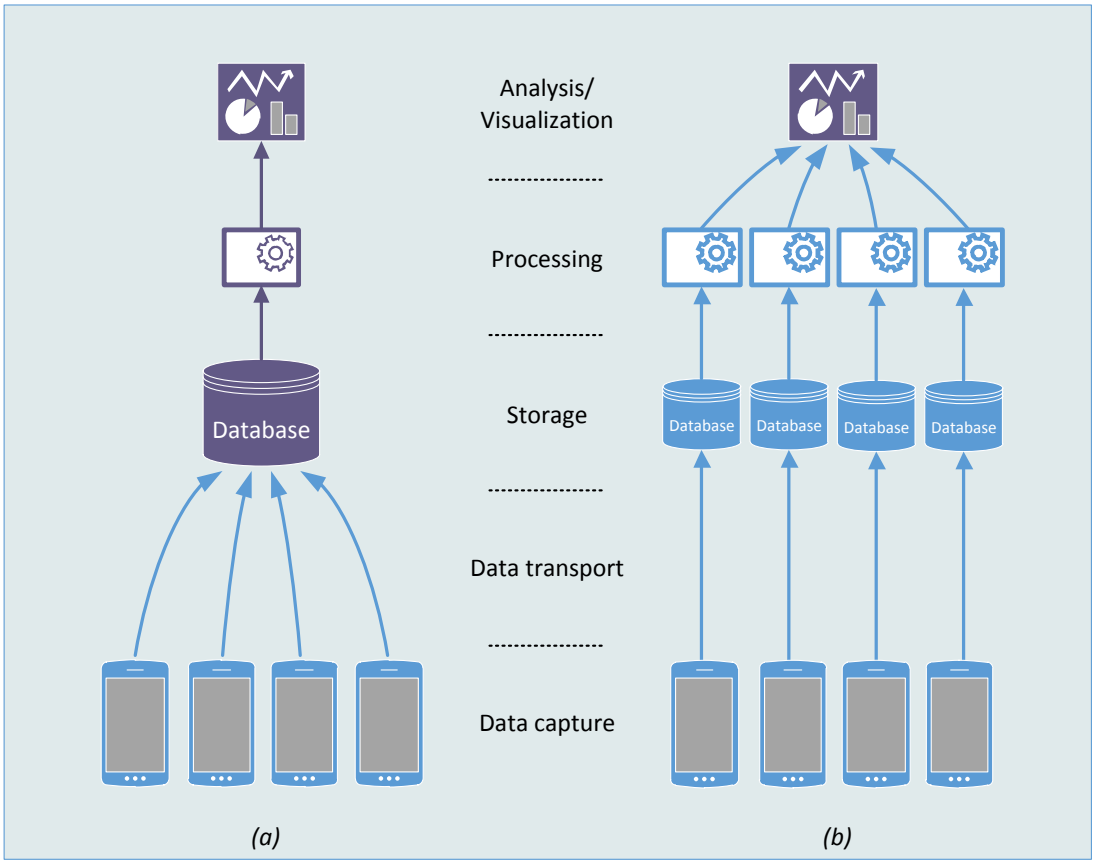

Fig. 1. Two different approaches in collecting and sharing personal data streams. In (a), data is collected in centralized database controlled by one entity, which then processes and analyses the data. In (b), storage and processing of data is decoupled from sharing and the owners of the data have better control over it.

lution is a critical asset in understanding large-scale, system-level interactions of technological and social phenomena.

Moreover, the privacy enhancing building blocks available today need to be integrated into an overall privacy engineering environment and to be deployed on servers, network infrastructures and devices, where they will interact with multiple other independently developed technologies. The way in which the privacy properties of these modules interact with each other and with the surrounding system through the entire protocol stack is not well understood. For example, interaction of different systems can lead to surprising effects (e.g., unwanted data flows) resulting in the violation of privacy policies or assumptions implicit in privacy technologies.

2) Tool Clinics: Another problem related to the lack of wide PETs adoption is the tendency of engineers designing those tools to focus on the solution of a problem as they perceive it, without adequate consideration of the viewpoints of other involved actors in the system, who might have different viewpoints and interests. One example is, of course, the end-user of a technological tool, who has different expectations and assumptions about the technology than the providing organization. If these assumptions and expectations are violated, the user is likely to reject the technology. Requirements from other stakeholders are also to be considered, like the interests and technical capabilities of those who host and manage the technology, or the role of regulators, etc.

The potential problems that may result from privacy researchers and practitioners working mainly in isolation from other stakeholders is extensively analyzed and discussed in the Dagstuhl Seminar report [5]. The authors identify the tendency to focus narrowly on only a single or maybe a few variables as part of the problem. The report encourages using methods for collectively reflecting on technology solutions and techniques, in order to critically assess their design from multiple perspectives. More specifically, it suggests introducing "tool clinics", meaning frameworks that encourage collaborative reflection of technological solutions or other artifacts, into research projects, courses, or conferences.

A tool clinic can be used to provide a setting for those who are developing the solutions to rethink the framing and presentation of their solutions [5]. So after one has identified the affordances of a technological solution and its possible consequences to people and society, one would have to collect the perspectives and practices from different experts, disciplines and stakeholders associated with it. The objective is to reflect from different perspectives on practices around the development, encoding, use, domestication, decoding and sustainability of a tool to gain quasi-ecological validation. The

challenge here would be to bring the various stakeholders together under the right format. This could be, for example, a session at an academic conference.

Definitely, one of the main stakeholders in this discussion is the user and in the next section it will become more obvious how things change, when one takes the perspective of users and consider their behavior regarding the protection of their online privacy.

III. Sharing Practices

"Sharing" is the central activity of mobile crowdsensing applications. It is thus a central requirement of any privacy and trust tool to "fit" the user's actual sharing practices and needs. Two main research thrust can be identified: behavioral privacy and consequence-based protection approaches.

A. Behavioral Privacy

Recent studies have shown that people frequently report preferences that they don't act upon in practice, and that their privacy concerns are highly sensitive to contextual factors [8]. Indeed, from an individual user's point of view, privacy has long since been reported not to be binary (i.e., not a simple yes/no decision) – rather, the decision of whether a particular piece of information is private or not depends on who is using the information (information receiver), and for what purposes (information usage). In fact, privacy can be seen as a continuum between confidentiality and disclosure, used in such a way to allow people to present themselves to their social surrounding appropriately according to the situation [9].

However, today's social networking tools are more often than not modeled around two distinct sharing states: "limited" and "public". Limited sharing applies to relatively static relationships with other members (e.g., friends, family), while public sharing with everybody usually applies to information that is congruent with the sharer's existing or desired public image. Users typically have to consciously decide which group a certain information can be shared with, in order to "fit" the desired proper role that should be presented. In ubiquitous experience sharing, the continuity of sensor streams makes such manual classifications increasingly hard.

How are owners of mobile devices going to decide whether to share their a particular sensory stream, and up to which point, or at which granularity? We see two broad factors that would influence people's disclosure preferences and decisions:

- The first factor is the degree of *trust* in the information receiver. This trust decision is influenced by past experience – or, in the absence of it – the degree of trust that other community members have placed in the information receiver and his intentions in how to use the information. To allow users to assess the information receiver, we need to *identify reliable trust signaling mechanisms* for the sensing context. Information sensed from the physical environment actually offers the opportunity to pick up several implicit trust signals (e.g., the user really is in the location she claims she is) and infer additional information about the behavior of the information receivers.
- The second factor is the *risk/benefit trade-off* associated with information usage, as perceived by the individual. With the risk of reducing one's privacy, this assessment will depend on the perceived value of incentives, e.g., monetary incentives for higher disclosure (such as the frequency, granularity and accuracy of data transmitted), and the social context (the reputation of the individual and the information receiver in community). Given the importance of reputation, we need to design systems that collect, aggregate and transmit all relevant reputation information accurately and reliably. This can be achieved by maximizing incentives and minimizing workload for users, but we need to *identify the specific incentives and workload perceptions for the sensing context*, and provide configuration tools that allow users to express their preferences correctly and efficiently.

B. Consequence-based Protection

For years, privacy protection tools have struggled to gain widespread adoption. For example, most modern Web browsers today support a variety of cookie control tools (e.g., the Ghostery[1] plug-in), yet many users still prefer to simply accept all cookies. At the same time, those who do use privacy tools are often unsure about the actual effects and benefits provided, e.g., as in the case of Facebook's bewildering privacy settings [10]. Moreover, as privacy tools are moving beyond simple On-Off dichotomies and into fine-grained preferences and policies, users are increasingly burdened with having to "debug" the consequences of their settings.

One major reason for this is often the technical, object-based orientation of such tools. Instead of setting access control rules on classes of objects (e.g., photos, "wall entries", comments sections), users often require a "translation" of the implication of such settings. For example, in order to understand the effects that allowing others to "tag" oneself in a photograph has, users need to understand that their personal pictures can now be found using a simple Web query. This problem becomes more and more challenging as information is combined into systems, with often complex implications that would need to be explained in clear language to the user.

A core requirement for novel privacy and trust interfaces in ubiquitous experience sharing systems is thus the need for visualizing or otherwise communicating the *consequences* of various privacy choices, rather than simply framing these at a technical level (i.e., "Feature: On/Off"). This would have implications on multiple levels: on a user interface level, it would lower the cognitive distance between the system image and a user's mental model; on a legal level, it would raise the quality of informed consent gathered by the data collection system; on a trust level, it would lower the potential for misunderstanding between the data controller and their "customer", thus raising consumer satisfaction.

Three core challenges need to be addressed in order to move privacy and trust interfaces from a system-level to a "consequence-based" level:

[1] See www.ghostery.com

1) *Expression of potential consequences:* The most challenging aspect of consequence-based privacy and trust tools is finding the right "language" to express the consequences of their choices to users. Too short descriptions run the risk of oversimplification, too long descriptions might border on "legalese" that become tedious to parse. Work on "privacy labels" [11] offers a first approach to making consequences of data collections more human-readable, though so far mostly in the domain of Web privacy. Another approach might be the use of simple metrics that can illustrate different effects, e.g., the number of people being able to read one's post or access one's current location.
2) *Decision support:* Just as people appreciate shopping recommendations and reviews from friends and family more than from random strangers, privacy choices could be "crowd-sourced" from a user's trusted contacts. By displaying actually choices from known contacts, potentially combined with expert advice that further explains these choices, users could receive substantial help and advice when having to make privacy-related choices. However, social compliance effects might turn out to force people into oversharing, so more research is needed in order to determine the effects of different user groups, different feedback, and different external information sources on actual privacy choices.
3) *Minimal effort:* Not all users will want to spend a significant amount of time adjusting their privacy preferences. Having different levels of detail for different sets of users and different situations can help to prevent overwhelming the individual with questions and decisions that might make users refrain from actively controlling their data sharing settings. Research in understanding the right moment and the right level of detail at which to solicit user preferences for data sharing, as well as the periodicity with which to re-confirm initial choices made earlier, is critical for building "minimal-effort" user interfaces that grow in complexity as users move from casual tool use to power users.

IV. Social Justice

Open challenges such as the market-dependent features of the PDS, the social and contextual nature of trust in sensing data sharing, and difficulties of minimal-effort decision support, all demonstrate that the challenges of participatory sensing reach beyond privacy. These challenges signify concerns about the impact of sensing on fairness and "social justice", i.e., the ability for people to flourish regardless of their social position [12]. During our Dagstuhl seminar it was discussed how to operationalize concepts like fairness, human flourishing, structural change, and balances of power [12] for the design of data collection features, processing, sharing, and user interfaces. One outcome was an analysis tool – a social justice impact assessment – to help system designers consider the social justice implications of their work during the design phase [5].

Ubiquitous sensing and information sharing may impact justice and fairness in a number of social domains. It seems clear that participatory sensing will have a positive impact on health, as patients are empowered with respect to institutions and new forms of data enable new kinds of diagnosis, monitoring and treatment. Similarly, care and independent living support for older citizens may be positively impacted by increased experience sharing. On a social level, we think ubiquitous sharing can benefit community integration through applications like neighbor-to-neighbor sharing of goods, services, and experiences; projects on walkability and bikeability in urban areas; and helping communities make the case about inequalities such as air pollution levels in underserved communities [13]. Applications such as poll watching and cop watching can increase transparency and accountability of powerful organizations to the citizen [14].

But there are also areas where participatory sensing might aggravate existing social problems and disparities. Digital vigilantism may be used to police social norms and reduce individual autonomy. Participatory sensing may allow for increased monitoring and measurement in schools, further quantifying student learning outcomes. Expanded tracking in criminal justice may impact the concept of rehabilitation and a "second chance." Participatory sensing also raises the specter of increasing inequality in social sectors that involve profiling and demographic sorting. With increased sensing capabilities, will we see the emergence of new indicators for marginalization? For example, the current trend in insurance is towards profiling to quantify individual risks. Participatory sensing data ranging from driving habits to location-based indexing of environmental data could all increase the granularity of personal profiles. Similar sorting could impact the financial industry, risk management, and price discrimination. New categories may be perceived as unfair if they are difficult to understand. Indeed, the complexity of the algorithms used to sort our individual "big data" may be quite difficult to explain and understand.

Because participatory sensing is likely to impact fairness and social justice in complex ways, we outlined a method for examining specific applications to evaluate their potential impact on social justice. An evaluation would take into account the stakeholders collecting, analyzing, and benefitting from sensing data. Building on earlier privacy assessment techniques [15], we outlined a procedure to break down the large issue of "social justice" into smaller component parts, and then examine how each of these smaller concepts might be impacted by an application, defining threats and mitigation strategies along the way. The detailed understanding of threats to social justice concepts, and ways to mitigate those threats through design or policy, can aid in the application design or regulation process.

Framing a social justice agenda for participatory sensing suggests many open questions for future research. For example, how do we encourage explicit social justice goals as part of application design? How can we integrate social justice impact assessments into ongoing technical work? As

participatory sensing projects become more common, research into the social impacts of sensing is also needed. For example, how will new forms of transparency enabled by participatory sensing impact individuals, powerful people, and institutions differently? Do new sensing technologies favor individual liberties at the expense of social action, or vice versa? And finally, research on the social impacts of big data will help us understand the fairness and justice implications of participatory sensing. For example, what factors in sorting and categorization processes help people feel that resulting algorithmic treatment is fair or unfair? New research from design to deployment to use of participatory sensing systems is needed to move beyond privacy towards understanding the social justice implications of participatory sensing [5].

V. Conclusions

Mobile crowdsensing – the act of ubiquitously sharing our sensory experiences with others – offers a tremendous potential for improving our lives, both at an individual level (e.g., personal health and wellbeing) as well as on a societal level (e.g., environmental monitoring, public health). However, as with many technological advances in the past, this carries a significant impact on personal privacy and individual self-determination. In order to avoid a dystopian vision of disadvantageous discriminations, unwanted advertisements, and identity theft, we must intensify our efforts to build privacy-respecting technologies that on one hand protect and guarantee the democratic value of informational self-determination and on the other hand are usable and adopted by people. This article has outlined several challenges across three distinct domains – privacy engineering, sharing practices, and social justice – in order to advance future research in this area.

VI. Acknowledgements

The contributions of the Dagstuhl seminar participants, as manifested in the corresponding report [5], provided the basis for this article. We are grateful for the wealth of ideas and insights from all participants and would like to explicitly acknowledge their contributions to the various topics of this article: The "Personal Data Services" group consisted of Alessandro Acquisti, Claudio Bettini, Rainer Böhme, Claude Castelluccia, Tassos Dimitriou, Frank Dürr, Raghu K. Ganti, Jens Grossklags, Deborah Estrin, Michael Friedewald, Renè Mayrhofer, David Phillips, Kai Rannenberg, Norman Sadeh, Marcello Scipioni. The "Tool Clinics" group members were Anthony Morton, Bettina Berendt, Seda Gürses, Jo Pierson. The topic of "Consequence-based Privacy Decision-making" was discussed by Zinaida Benenson, Delphine Christin, Alexander De Luca, Simone Fischer-Hübner, Thomas Heimann, Joachim Meyer. The issue of "Social Justice" was explored by Mads S. Andersen, Ian Brown, Ioannis Krontiris, Sören Preibusch, Angela Sasse, Katie Shilton, and Sarah Spiekermann. We would also like to thank Alessandro Acquisti for valuable feedback on earlier drafts of this article.

References

[1] R. K. Ganti, F. Ye, and H. Lei, "Mobile Crowdsensing: Current State and Future Challenges," *IEEE Communications Magazine*, vol. 49, no. 11, pp. 32–39, 2011.

[2] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick, "A survey on privacy in mobile participatory sensing applications," *Journal of Systems and Software*, vol. 84, no. 11, pp. 1928–1946, Nov. 2011.

[3] G. Danezis and S. Gürses, "A critical review of 10 years of privacy technology," in *Proceedings of Surveillance Cultures: A Global Surveillance Society?*, London, UK, April 2010.

[4] S. Fischer-Hübner, C. Hoofnagle, I. Krontiris, K. Rannenberg, and M. Waidner, "Online Privacy: Towards Informational Self-Determination on the Internet," *Dagstuhl Manifestos*, vol. 1, no. 1, pp. 1 – 20, 2011.

[5] A. Acquisti, I. Krontiris, M. Langheinrich, and M. A. Sasse, "'My Life, Shared' - Trust and Privacy in the Age of Ubiquitous Experience Sharing (Dagstuhl Seminar 13312)," *Dagstuhl Reports*, vol. 3, no. 7, pp. 74–107, 2013.

[6] R. Cáceres, L. Cox, H. Lim, A. Shakimov, and A. Varshavsky, "Virtual individual servers as privacy-preserving proxies for mobile devices," in *Proceedings of the 1st ACM Workshop on Networking, Systems, and Applications for Mobile Handhelds*, ser. MobiHeld '09, Barcelona, Spain, 2009, pp. 37–42.

[7] M. Y. Mun, D. H. Kim, K. Shilton, D. Estrin, M. Hansen, and R. Govindan, "PDVLoc: a personal data vault for controlled location data sharing (In press)," *ACM Transactions on Sensor Networks*, 2014.

[8] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *IEEE Security and Privacy*, vol. 3, no. 1, pp. 26–33, Jan. 2005.

[9] M. Raento and A. Oulasvirta, "Designing for privacy and self-presentation in social awareness," *Personal and Ubiquitous Computing*, vol. 12, no. 7, pp. 527–542, Oct. 2008.

[10] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing facebook privacy settings: User expectations vs. reality," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 61–70.

[11] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder, "A nutrition label for privacy," in *Proceedings of the 5th Symposium on Usable Privacy and Security*. ACM, 2009, p. 4.

[12] A. Sen, *The idea of justice*. Cambridge, MA: Harvard University Press, 2009.

[13] P. Dutta, P. M. Aoki, N. Kumar, A. Mainwaring, C. Myers, W. Willett, and A. Woodruff, "Common Sense: Participatory Urban Sensing Using a Network of Handheld Air Quality Monitors," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys '09)*, New York, NY, USA, 2009, pp. 349–350.

[14] L. Hueya, K. Walby, and A. Doyle, *Surveillance and security: technological politics and power in everyday life*. Routledge, 2006, ch. Cop watching in the downtown eastside: exploring the use of (counter) surveillance as a tool of resistance, pp. 149–165.

[15] M. C. Oetzel and S. Spiekermann, "Systematic methodology for privacy impact assessments," *European Journal of Information Systems*, July 2013.