

Exploring Privacy Concerns about Personal Sensing

Predrag Klasnja¹, Sunny Consolvo², Tanzeem Choudhury³, Richard Beckwith⁴,
and Jeffrey Hightower²

¹The Information School, University of Washington, Seattle, WA, USA

²Intel Research Seattle, Seattle, WA, USA

³Department of Computer Science, Dartmouth College, Hanover, NH, USA

⁴Intel Research, Portland, OR, USA

More and more personal devices such as mobile phones and multimedia players use embedded sensing. This means that people are wearing and carrying devices capable of sensing details about them such as their activity, location, and environment. In this paper, we explore privacy concerns about such *personal sensing* through interviews with 24 participants who took part in a three month study that used personal sensing to detect their physical activities. Our results show that concerns often depended on *what* was being recorded, the *context* in which participants worked and lived and thus would be sensed, and the *value* they perceived would be provided. We suggest ways in which personal sensing can be made more privacy-sensitive to address these concerns.

Introduction

Personal devices with embedded sensing are becoming pervasive. GPS units are present even in midrange mobile phones, and with the popularity of the iPhone and similar multimedia-oriented devices, accelerometers and proximity sensors are quickly moving into the consumer mainstream. Incorporating sensing into such personal, mobile devices enables a range of compelling applications, from location-based services—getting a restaurant recommendation near the user’s current position, for example—to real-time detection of the user’s physical activities.

However, having sensors embedded in devices that users wear or carry with them all day, every day can also be problematic. For example, having one’s location constantly sensed can enable an unwanted person to learn where and when a user spends her time. Such information could potentially enable stalking or other types of criminal activity. In addition to such security considerations, people may simply be uncomfortable with others knowing their location, or even with their location being sensed in the first place. Mobile applications such as the Audio Loop [3], which continuously record raw audio, also raise concerns and introduce issues around how (or even if) to obtain consent to be recorded from others whose data might be captured by the user’s device [5]. Such concerns could affect the adoption and use of devices that embed sensing and introduce problems into social relationships.

The usefulness of continuous sensing to enable a wide range of applications on the one hand and the potential privacy and security issues that accompany sensing on the

other, raise a design challenge for pervasive computing researchers. That is, how can we design such systems so that they use sensing when and where it is needed while respecting the privacy and comfort of users and others who may be monitored?

In this paper, we address this question by describing results from a study where 24 participants who used their mobile phone and a personal sensing device to track their physical activities for three months were interviewed about their reactions to and speculations about personal sensing in everyday contexts. As part of the study's exit interviews, participants were asked about any concerns they had with the sensing that was employed during the study and to speculate about other sensors that could be added to improve the system's activity inference capabilities. Our results reveal that privacy concerns varied greatly depending on *what* the sensor was recording, the *context* in which participants worked and lived and thus would be sensed, and the *value* they perceived in the capabilities that would be enabled. Participants often weighed the intrusiveness of what was being monitored about them and potentially others with whom they come in contact against perceived application benefits. If the latter were not seen as compelling enough, users would reject the use of sensing.

In what follows, we describe our method, results, and discuss implications of the results for the design of everyday personal sensing technologies. We then review related work and conclude.

Method

The data presented in this paper come from the exit interviews from a three-month field study of the UbiFit system which used mobile technology to encourage physical activity. Results that focused on how the system affected awareness and behavior related to physical activity and the effectiveness of the persuasive elements of the system have been presented elsewhere [2]. In this paper, we present results on participants' reactions to the sensors that were used during the study and other sensors that we were considering for a future version of the system to improve its activity inference capabilities. We describe the participants and our interview method below.

Twenty-eight people (15 female, aged 25-54), recruited from the Seattle metropolitan area's general public by a market research agency, participated in our three-month field study of the UbiFit system. The participants represented a range of professions, including real estate agent, personal care assistant, psychologist, teacher, comedian, public relations specialist and others. Twenty-four (14 female) participants took part in the sensing portion of the exit interview. The sensing questions were not asked of the remaining four because of time constraints.

During the study, 15 of the 24 participants wore a fitness device that used a 3-D accelerometer and a barometer to automatically infer walking, running, cycling, stair machine, and elliptical trainer activities. The remaining nine participants manually kept track of their physical activities using a journal on their mobile phones; participants who wore the fitness device also used the phone journal to record activities that the device was not trained to detect (e.g., swimming, yoga). Toward the end of the exit interview, we asked participants how they thought the fitness device inferred activities, and then explained to them how it did so. As part of our

Exploring Privacy Concerns about Personal Sensing

explanation, we provided a printout of the accelerometer and barometer data for three types of activities—sitting, walking, and running—to illustrate how the device could distinguish different activities. We then asked if they had any concerns about this data being recorded about them all day, every day and stored on their phones and, potentially, on a companion web site (which was not part of the study, but was something we were considering adding in future work).

We then suggested potential improvements, including providing flexibility on where the device could be worn (for the study, it had to be worn on the waistband), improving the accuracy of activity inference, inferring more information about activities, and inferring more types of activities. After getting participants' feedback about the usefulness of the suggested improvements, we explained that to implement the improvements would require the use of additional sensors, and we wanted to get their reactions to two sensors we were considering: (1) GPS and (2) a microphone. For each, we asked if the participants were familiar with the sensing technology and explained what it recorded and how it could be used to make the improvements. We also showed the participants a map of a run, along with the distance, elevation, and pace information for the run that was derived from GPS data. We then asked the participants to speculate on how they would feel about the sensors running all day, every day on their device, as the accelerometer and barometer had done in the study.

After they shared their initial thoughts, we probed about the implications of using the sensors. To ensure that the trade-offs were fully weighed, for people who were positive about the sensors, we brought their attention to possible concerns, and for people who were concerned, we pointed out the benefits that the sensors would enable. In addition, we suggested that the system would not have to keep the raw sensor data indefinitely, but could instead only maintain a small window of sensor data needed to calculate higher level measures such as distance and pace, after which the sensor data could be discarded. Finally, for audio, we noted that we might not need raw audio, but could record only certain frequencies within the audio stream that were needed for the inference. Such filtered audio would not contain enough information to ever reconstruct the content of a conversation, although an audio expert still might be able to determine that a conversation took place, how many people participated, what their genders were, and the general emotional tone of the conversation (e.g., whether it was an argument, a happy conversation, and so on). To clarify this idea, we played two audio recordings of the same nine second clip of a conversation between two males. One clip was raw audio, recorded by a microphone that was worn the same way the fitness device was worn for the study (i.e., clipped to the waistband). The other clip was the same recording, filtered to remove the unnecessary frequencies as described above [9]. We then asked participants how they felt about the filtered audio versus the raw audio and if and for how long they would be comfortable keeping that type of audio recording.

Interviews were audio-recorded and transcribed. The interview data was analyzed using open coding, a standard method of analyzing qualitative data.

Results

In what follows, we discuss how participants' reactions varied for different types of sensors, their speculations on data retention, the context in which the participants were likely to use the system, and the value that the participants saw in the functionality that the use of the sensing would enable.

Reactions to Different Sensor Types

Not surprisingly, participants reacted differently to the different types of sensors. None of the participants had any concerns about the accelerometer and barometer—the two sensors that were used during the field study. The participants did not consider this data to be particularly sensitive, and therefore had no problem with these sensors recording all day, every day, and for the data to be stored indefinitely on their mobile phone or on the fitness device. In addition, all participants who expressed wanting a companion fitness web site had no problem with the raw accelerometer and barometer data being stored there as well.

Reactions to GPS were more mixed. Unlike the accelerometer or barometer data, participants tended to think of the GPS data as being sensitive. 42% (14 of 24) of the participants had concerns about GPS being recorded all-day, every day. Concerns ranged from physical security—someone might get hold of the data and be able to figure out where the participants live or where their children go to school—to simply thinking that it was “creepy” or “big brother”-like. One participant commented that he does not like technologies that “*show where a human is exactly*” {Participant P3} and another one commented, “*I don't know about that...it can tell where you live and...that might be a little bit too much*” {P22}. When asked why GPS was different in terms of “being tracked” than the accelerometer, one participant explained that with the accelerometer “*it's not as specific, so the accelerometer isn't going to say that she is at <the intersection of> First and Pine*” {P22}. Nine of the 14 (64%) participants who were open to having the GPS data remain on their mobile phones also did not mind having the data stored on a companion web site.

Reactions to the raw audio were nearly unanimously negative. Only two of the 24 participants (8.3%) would consider a microphone that continuously recorded raw audio. Other participants adamantly replied that they were not willing to be recorded all the time, that they were uncomfortable being watched all the time (being recorded felt “*Big Brother-ish*” {P25}, “*I think I would feel too watched and too listened to*” {P22}), and even if they did not have a problem with being continuously recorded, they did not feel that it was okay to record those with whom they came into contact.

Recording audio in just the frequencies needed for activity inference was more acceptable. 25% (6 of 24) of the participants were willing to use the filtered audio and keep the data on their phones indefinitely (although only three were willing to upload this data to a companion web site). However, most participants remained uncomfortable. They simply found any audio recording to be too intrusive. The two most frequent classes of concern were 1) it made some participants feel as though they were being watched, and 2) even the filtered audio was seen as containing data that was too sensitive. One participant commented that “*I mean, even filtered, just I*

Exploring Privacy Concerns about Personal Sensing

don't know. I would feel too exposed, but with this device [with just the accelerometer and barometer], I could care less" {P22}. Another commented that filtered recording "*still just has that Big Brother effect to it*" {P27}. Regarding the sensitivity of the information, one participant commented that someone could still determine if you were with someone else, and another that the number of people in the conversation and its emotional tone were still "*a substantial amount of information*" {P2}.

Data Retention

In some cases, concerns about seemingly invasive sensors could be mitigated by changing the length of time that data were retained. While nearly half of the participants were unwilling to use GPS if the raw data (e.g., the latitude and longitude coordinates) were kept, all but one participant were willing to use it if the raw data were kept only for as long as was necessary to calculate the characteristics of detected physical activities (e.g., distance or pace of a run), and then promptly discarded. The exact length of the data window that the participants thought was acceptable varied, but most who wanted data purging thought that retaining one to 10 minutes of raw data at a time, unless a physical activity is being detected, was reasonable.

We found similar results for audio. A sliding data window of no more than one minute at a time of raw audio data was acceptable to 29% (7 of 24) of participants, although the majority (71%) found recording of any raw audio too invasive. Filtered audio fared better, however. If only a 10 minute sliding window of filtered audio was being saved, except for times when a physical activity is being detected, 62.5% (15 of 24) of participants were willing to use the microphone to get better activity detection.

Influence of Context of Use on Sensor Acceptability

Participants who worked in environments that required confidentiality unanimously objected to all forms of audio recording. When the use of a microphone was brought up, one participant explained that "*at work I'm privy to sensitive information that other people can't hear*" {P1}. Having a recording device of any sort was seen as completely unacceptable in that context, even if the audio was being continuously purged, often for fear that the device would be somehow compromised. Another participant, a psychologist, said that even the filtered audio contained too much information to be acceptable to use in patient consultations. She was afraid that such recordings could be subpoenaed in a potential legal case involving a patient, and that the risk was just not worth it. Even one of the two participants who would have been open to raw audio recording realized that his place of work would not have been okay with it, making the use of raw audio recording untenable for him.

The acceptability of using and storing GPS data was judged in a similar way. One participant commented that having data that shows where one is going makes her "*super uncomfortable*" since she has had friends with "*really controlling husbands,*" who would abuse such information {P3}. Another worked at "*a confidential site*" {P15}, and was concerned that someone could get access to the location data that he deemed confidential. GPS was generally more problematic for women who tended to

feel more vulnerable than the men did. The characteristics of the context in which a sensing system would be used, such as the confidentiality requirements of a workplace, or the perceived vulnerability of the user—strongly influenced how the sensing technology was judged.

Value of Sensing-enabled Applications

How much value participants perceive the data would provide was a factor in how they evaluated the acceptability of different types of sensors and data management. Runners who wanted to have maps of their running routes so they could plan future workouts were more willing to retain raw GPS data than were runners who did not think that those maps were particularly valuable. The latter group wanted to calculate the higher level information about their runs—the pace, distance, and elevation changes—but preferred not to keep the location information. Similarly, the psychologist reasoned that while she would really like the fitness device to automatically detect more types of activities, the risk and discomfort that any form of audio recording brought up for her far outweighed the benefit. She preferred to keep the device as it was and to continue to journal anything else manually. Another participant, who saw an additional benefit in the filtered audio (i.e., using the data to get feedback about her emotional reactions in different social situations such as being on a date), ultimately determined that she would not use it as she would feel obligated to explain to the people with whom she interacted that and how they were being recorded, and that would have been too complicated. Even though she found this emotional feedback idea really appealing, its value was not high enough to justify the potential harm it could do to her personal relationships.

The usefulness of a map of a run, the decrease in burden by having activities inferred automatically, and the value of other anticipated applications of the data such as emotional feedback, were weighed against other factors—legality, intrusiveness, and social etiquette—to determine whether the form of sensing needed to enable the desired functionality was deemed acceptable. If the benefits were not seen to clearly outweigh the perceived costs, the sensing method was rejected.

Discussion

The acceptability of personal sensing is a result of making trade-offs between the perceived value of an application and the costs—legal, social, and psychological—that the user perceives given the context in which she lives and works. While some researchers have argued that over time, changes in legal policy and social contract will decrease privacy concerns (e.g., [3]), it is unknown how far-reaching such changes are likely to be. For instance, many of the privacy issues raised in Warren and Brandeis's classic 1890 paper [8] remain relevant today. Moreover, there are clearly situations—such as attorney-client and doctor-patient interactions—where the need for confidentiality and privacy will not go away. Enabling users to make privacy trade-offs in an informed, educated way will be a key task for designers of sensor-enabled personal devices.

Exploring Privacy Concerns about Personal Sensing

Our results suggest at least three ways in which the acceptability of sensing can be increased, while respecting privacy. First, sensor data should be saved only when relevant activities are taking place. Results for both GPS and audio revealed that continuously purging the raw data increased user acceptance of both sensors. Second, whenever possible, a system's core functionality should be based on minimally-invasive sensing. The users can then be given a choice to decide whether to enable additional functionality that might require more invasive sensors. Physical activity detection, much of which can be done with a simple 3-D accelerometer, is a good example of a domain where such graded sensing could be implemented. And third, researchers should explore ways to capture only those features of the sensor data that are truly necessary for a given application. This means, however, that sensor systems might need to have enough computational power to perform onboard processing so that each application that uses a sensor can capture only the information that it needs.

We also note that users can make informed privacy trade-offs only if they understand what the technology is doing, why, and what the potential privacy and security implications are. Building visibility into systems so that users can see and control what data is being recorded and for how long supports informed use. Determining how this can best be done is a difficult, but important, design challenge.

Related Work

In their work on the Audio Loop, a memory aid that continuously records a sliding window of 15 minutes of raw audio, Hayes et al [3] found that while over half of the participants in their lab study raised privacy concerns, the four participants in a field study were positive about the system. At least two reasons could explain why our data indicate lower acceptance of raw audio recording than Hayes et al found in their field study. First, the systems are different, and participants may have therefore valued their functionality differently. While the benefit of improving the physical activity inference in our system often did not warrant the use of invasive sensing, the Audio Loop's functionality might have been perceived as being valuable enough. Additionally, the core functionality of our system used sensors that did not raise privacy concerns; rather, the more invasive sensing would have been used to improve the functionality. Second, it is unclear if Hayes et al's participants encountered the types of confidential situations that were common for participants in our study. This potential difference in the context of use—what Hayes et al [4] call *social knowledge* affecting privacy perceptions—might explain why our results are different.

Iachello & Abowd's [6] *proportionality method* offers a design method aimed at ensuring that privacy is taken into account throughout the design process. Our data support their emphasis on *desirability*—making sure that the system's value makes any privacy compromises acceptable—and offer specific ways (e.g., graded sensing and sensor data filtering) to achieve design *appropriateness* and *adequacy*.

Nguyen et al [7] looked at people's responses to everyday tracking and recording technologies (TRTs), and found that people were highly concerned about privacy issues in the abstract but were simultaneously unconcerned with the TRTs of everyday life. They argue that familiar technologies provide known benefits and that

risk is more abstract, rarely having been experienced. Their findings suggest that for users who already carry the sensors on a device they own (e.g., a GPS enabled phone), they may be more willing to adopt the services. Hayes et al [4] emphasize the role of users' experiences in shaping privacy perceptions of new pervasive systems.

Finally, Beckwith & Mainwaring [1] found that users' privacy concerns depend on their understanding of the technology they are using. Making good privacy decisions is difficult if the technologies are poorly understood. The high level of concern with GPS and audio that we found in this study is likely due to the higher level of understanding that the participants had about these technologies.

Conclusion

This paper examined user reactions to four different types of sensors—accelerometer, barometer, GPS, and microphone—that can be used to infer physical activities. The reactions were obtained in interviews after 24 participants had used a mobile phone and/or personal sensing device to track their physical activity for three months, grounding their reactions and speculations in real world use. We found that *what* data is sensed and recorded, the *context* in which the sensing takes place, and the perceived *value* provided by the sensed data influenced the privacy trade-offs participants were willing to make. We suggest that conservative recording and data retention policies, graded functionality, and giving users visibility and control over which sensors are used could help with the adoption of systems that use continuous personal sensing.

References

1. Beckwith, R. and Mainwaring, S. Privacy: Personal information, threats, and technologies. *Proc. ISTAS '05*, IEEE (2005), 9-16.
2. Consolvo, S., Klasnja, P., McDonald, D.W., Avrahami, D., Froehlich, J., LeGrand, L., Libby, R., Mosher, K. and Landay, J. Flowers or robot armies? Encouraging awareness & activity with personal, mobile displays. *Proc. UbiComp '08*, ACM Press, (2008), 54-63.
3. Hayes, G.R., Patel, S.N., Truong, K.N., Iachello, G., Kientz, J.A., Farmer, R. Abowd, G.D. The Personal Audio Loop: Designing a Ubiquitous Audio-Based Memory Aid. *Proc. Mobile HCI 2004*, LNCS 3160, Springer Verlag (2004), 168–179.
4. Hayes, G.R., Poole, E.S., Iachello, G., Patel, S.N., Grimes, A., Abowd, G., and Truong, K.N. Physical, social, and experiential knowledge in pervasive computing environments. *IEEE Pervasive Computing*, 6, 4, 56-63, 2007.
5. Iachello, G., Truong, K. N., Abowd, G. D., Hayes, G. R., and Stevens, M. Prototyping and sampling experience to evaluate ubiquitous computing privacy in the real world. *Proc. CHI '06*, ACM Press (2006), 1009-1018.
6. Iachello, G. and Abowd, G. From privacy methods to a privacy toolbox: Evaluation shows that heuristics are complementary. *ACM Transactions of CHI*, 15, 2, 8:1-8:30, 2008.
7. Nguyen, D., Kobsa, A., and Hayes, G. An Empirical Investigation of Concerns of Everyday Tracking and Recording Technologies. *Proc. UbiComp '08*, ACM Press (2008), 182-191.
8. Warren, S.D. and Brandeis, L.D. The right to privacy. *Harvard Law Review*, 4, 5, 1890.
9. Wyatt, D., Choudhury, T., and Kautz, H. Capturing spontaneous conversation and social dynamics: A privacy sensitive data collection effort. *Proc. ICASSP*, 2007.