

Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment

Andrew Raij
University of South Florida
raij@usf.edu

Santosh Kumar
University of Memphis
santosh.kumar@memphis.edu

Animikh Ghosh
SETLabs, Infosys
Animikh.Ghosh@infosys.com

Mani Srivastava
UCLA
mbs@ucla.edu

ABSTRACT

Wearable sensors are revolutionizing healthcare and science by enabling capture of physiological, psychological, and behavioral measurements in natural environments. However, these seemingly innocuous measurements can be used to infer potentially private behaviors such as stress, conversation, smoking, drinking, illicit drug usage, and others. We conducted a study to assess how concerned people are about disclosure of a variety of behaviors and contexts that are embedded in wearable sensor data. Our results show participants are most concerned about disclosures of conversation episodes and stress — inferences that are not yet widely publicized. These concerns are mediated by temporal and physical context associated with the data and the participant's personal stake in the data. Our results provide key guidance on the extent to which people understand the potential for harm and data characteristics researchers should focus on to reduce the perceived harm from such datasets.

Author Keywords

Privacy, Information Disclosure, Wearable Sensors, Mobile Health, User Study

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*evaluation/methodology, user-centered design*; D.2.1 Requirements: [elicitation methods]; K.4.1 Public Policy Issues: [Privacy]

General Terms

Experimentation, Measurement, Human Factors

INTRODUCTION

Networked body-worn sensors and those embedded in mobile devices we carry (e.g., smartphones) can collect a variety of measurements about physical and physiological states,

such as acceleration, respiration, and ECG. By applying sophisticated machine learning algorithms on these data, rich inferences can be made about the physiological, psychological, and behavioral states and activities of people. Example inferences include dietary habits, psychosocial stress, addictive behaviors (e.g., drinking), exposures to pollutants, social context, and movement patterns. These inferences, particularly when made continuously as people go about their daily lives, have many uses, such as sharing activity and context information with friends and family [35], self-monitoring health for behavior change [7], and scientific study of human physiology, psychology, and behavior [2, 10].

New Privacy Concerns and Challenges

While useful, such sensing systems raise significant new privacy concerns. Seemingly innocuous data shared for one purpose can be used to infer private activities and behaviors that the individual did not intend to share. For example, inertial sensor data (e.g., accelerometers and gyroscopes) shared with caregivers for evaluating and improving gait or physical activity levels could also be used to track location and movement patterns by sophisticated tracking algorithms [12] that make use of public information such as street maps and an individual's work locations. Inertial data may also reveal sensitive medical conditions, such as seizures [22], that one may wish to keep private. As another example, respiration and location data, combined with publicly available pollution maps, could be shared to measure an individual's exposure to air pollution, where respiration measurements are used to estimate the amount of air inhaled. However, this same data can also reveal the timing and duration of conversations [25] or even smoking [2]. In some cases, inferences from sensory data (e.g., public speaking inferred by fusing prolonged speaking episode and elevated stress) combined with other public information (e.g., department and university) can also re-identify an individual.

Privacy research has traditionally dealt with reidentification from quasi-identifiers [32], search histories [3], movie ratings [27], and social networks [28]. Some work exists on understanding the privacy concerns emerging from sensory data, such as location traces [16], but little work has investigated the new privacy concerns that emerge from the disclosure of measurements collected by wearable sensors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

While the privacy risks emerging from sharing data collected by personal sensors are poorly understood, their adoption is growing rapidly. One example of such a system is AutoSense, an experimental, unobtrusive wearable sensor suite capable of capturing physiological data and using it to infer the wearer's behavior and psychological state in real-time [2]. It can be worn for weeks at a time and is thus capable of collecting significant amounts of personal data from the mobile environment of its wearer. To date, AutoSense has been used by 60+ participants in behavioral science field studies totalling 1,000+ hours, where it collected ECG, respiration, accelerometer, temperature, and skin conductance data; and from this data made continuous inferences of physical activity, posture, stress, conversation, and commuting. This rich and unique dataset, however, can not be readily shared with other researchers out of fear of similar unknown threats to privacy. This is because seemingly innocuous personal datasets have previously been shared, only to learn later of hidden privacy threats within them [27, 13]. We note that commercially available commodity devices capture similar information [1, 36]. To see continued growth and adoption of such devices, the privacy challenges associated with them must be addressed.

Contributions

Via a user study, we study the privacy concerns associated with the disclosure of data collected by wearable sensors in the mobile environment. We first develop a conceptual framework for examining the privacy issues associated with the disclosure of continuously-collected physiological, psychological, and behavioral data. Next, we create a new privacy survey to assess how concerns about privacy threats change as various behaviors and contexts are restricted and abstracted. Lastly, to analyze how concern level changes as the respondents' stake in the data is increased [24], we administered the survey to two groups. The first group, Group *NS* ($N = 36$), completed the survey and had no personal stake in the data described. The second group, Group *S* ($N = 30$), was composed of participants in an AutoSense field study who had direct personal exposure to the collection of sensitive datasets. They wore the AutoSense system for 3 days. Group *S* completed the privacy survey twice, after the 3 day collection period, and again after they were visually presented with behavioral and contextual inferences derived from their data (e.g., periods of stress, conversation, commuting, etc.) to further increase their personal exposure to and overall stake in their data.

We analyze the data from three perspectives. First, we assess how disclosure of different behaviors and contexts affect participant concern levels as their stake is increased in the data. Second, we evaluate the impact of applying various restrictions and abstractions on concern level. Third, we assess the impact of reidentification on the concern level as the role of the data consumer (i.e., whom the data is disclosed to) is varied from the research team to the general public.

Key Results. We find that participants are most concerned about release of conversation and commuting behavior, and release of stress, a psychological state. However, those with

low or no personal stake in the data did not indicate nearly as strong concerns, implying that a person needs to have a personal connection to the data to understand the implications to his/her privacy. Restricting or abstracting the disclosure of temporal context (e.g., releasing duration in place of timestamps) had a significant effect on reducing privacy concerns for a variety of contexts and behaviors. Lastly, the risk of reidentification doubled concerns regarding disclosure of the data; the effect was largest for disclosure to the general public. Our results call for deeper investigation into the new privacy issues emerging in the domain of personal sensing.

RELATED WORK

Understanding and Awareness of Privacy Risks: Several well-publicized privacy breaches have contributed significantly to awareness of privacy risks [32, 3, 27, 28]. This awareness has led to studies that seek to understand the privacy concerns of users when sharing personal data, such as location [24, 33, 8, 4, 35, 5], calendars [31], and more generally sharing over online social networks [34]. However, these studies deal with sharing information whose privacy implications are already understood by most individuals (e.g. sensory information such as location and non-sensory information such as calendars). Their insights do not necessarily apply to the rich set of sensors in wearable devices and mobile phones that enable non-obvious inferences about users' behaviors and activities.

Some recent papers have examined awareness of non-obvious privacy threats from personal sensing applications and WiFi data. Using the Personal Audio Loop, Iachello et al examined privacy concerns regarding audio capture of conversation [15]. However, they focus on the privacy concerns of secondary stakeholders and third parties, who may be inadvertently recorded by the system. Our work focuses on the concerns of the primary user whose conversation episodes may be inferred without any recording of the audio, i.e., from an innocuous respiration sensor. Klasnja et al [17] investigate privacy concerns by interviewing participants using a physical fitness system. Concerns were evaluated for activity, GPS, and audio data. Our work is complementary, in that it studies concerns for physiological sensors and the psychological and behavioral inferences made from them. Klasnja et al [18] and Consolvo et al [6] investigate privacy problems associated with daily WiFi use. This work employs similar methods to assess privacy concerns, i.e., we also show participants their own potentially-sensitive data and then assess how exposure to the data changes awareness and concerns. However, we apply it to expose concerns associated with the use of innocuous physiological sensors.

To our knowledge, we are the first to highlight privacy concerns associated with behavioral and affective inferences that can be made from seemingly innocuous physiological data (e.g., ECG and respiration). Historically, these sensors have not been viewed as sources of privacy threats. We also identify disclosure of psychological state (e.g., stress) as concerning to users.

Maintaining Privacy: Increased awareness of privacy risks

has led to anonymization techniques [32, 23, 21] that seek to thwart reidentification attacks on personal data from which explicit identifiers have been removed. Reidentification attacks exploit quasi-identifiers such as age, zip code, etc., which cannot be entirely suppressed due to their utility to the end-user, but can reveal identity when combined with background information. Current techniques, mostly explored in context of relational data, aim to preserve an individual's identity in a population. They combat reidentification with techniques such as suppression, perturbation, and generalization. While analogous techniques have utility for sensor time series data [11, 26], anonymizing sensory information is often harder and different because sensor data (i) may be both sensitive and quasi-identifying, making it harder to achieve privacy without affecting utility, and (ii) may need to be shared with identity (e.g., with caregivers). Thus, existing anonymization techniques alone cannot be used to protect individuals sharing personal sensor data. New approaches are needed to preserve *behavioral privacy* when identity cannot be removed from the dataset.

REASONING ABOUT PERSONAL SENSING PRIVACY

This section describes a conceptual framework for examining the privacy issues associated with physiological, psychological, and behavioral data captured by personal sensors (see Figure 1). Our framework is primarily concerned with what [30] calls the *disclosure boundary*, i.e., where privacy and publicity are in tension. We focus specifically on the choices data producers can make that displace this disclosure boundary for personal sensory data. In addition, the threats highlighted by our framework touch upon the *identity boundary*, and how unintended disclosure of an individual's behavior may change how others view the individual.

Our framework is composed of six elements: measurements, behaviors, contexts, restrictions, abstractions, and privacy threats. Data producers capture *measurements* with sensors, which can then be processed to infer *behaviors* and *contexts*. To control access to the data, data producers can put *restrictions* in place that prevent sharing of specified measurements and/or behaviors and contexts derived from them. To control the level of detail in the shared data, *abstractions* can be applied to unrestricted measurements, behaviors, and contexts. The set of disclosed behaviors and contexts, and the level of abstraction applied to each, ultimately decide the types and significance of the resulting *privacy threats*.

We motivate each element of the framework using a fictional AutoSense study participant, Jane Doe. Jane is participating in a study to examine physiological, psychological, and behavioral factors associated with stress. Jane is nearing the end of the third of seven consecutive days wearing the AutoSense system. She had a stressful day at work, so she decided to run home for exercise and stress relief. Her husband, John Doe, met her in the middle of the run, and they completed the run together. Jane has a medical condition that causes occasional seizures, and after the run, Jane experienced a seizure. AutoSense recorded the following information about Jane's run, among others:

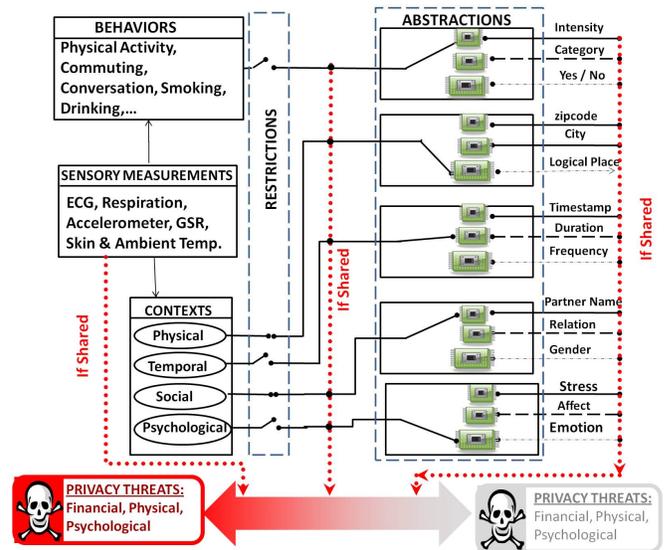


Figure 1. A conceptual framework for reasoning about privacy issues in sharing personal sensory data.

- Jane put on the sensors on September 21 at 8:00am.
- Additional processing of the accelerometer data indicated Jane was running September 21, from 5:33pm to 6:07pm, and she experienced a seizure from 6:15pm to 6:25pm.
- GPS data indicates the run started at 185 N Industry Street in Phoenix, AZ and ended at 1359 W Suburbs Street in Phoenix, AZ. The route taken was also recorded.
- Jane was angry and stressed from 5:25pm to 5:51pm. She was relaxed from 5:51pm until the start of her seizure at 6:15pm. She was terrified from the start of her seizure at 6:15pm until she removed the sensors at 6:30pm.
- Starting at 5:28pm, another mobile phone with the Bluetooth ID "JohnDoe" entered her phone's proximity and remained there until Jane removed the sensors at 6:30pm.

Measurements, Behaviors, and Contexts

Measurements are the raw data captured by wearable sensors, such as ECG, respiration, and accelerometer. **Behaviors** are actions taken or experienced by the data producer, and are inferred from measurements. In Jane's case, both the running episode and the seizure could be inferred from her accelerometer measurements. Jane was aware of the system's ability to detect running because it was mentioned in the study's informed consent document. However, the document did not mention that accelerometer data could also capture the motion signature of her seizure. The study designers anticipated behaviors that are not of interest might be captured by the sensors, but they did not anticipate that seizures, a sensitive medical condition, could be captured.

Other similar accidental captures of private behavior are possible. For example, heart rate and respiration, measurements used to infer stress level could also be used to infer cocaine, heroin, and other illicit drug use, all of which affect these measurements in extreme ways. Smoking, alcohol consumption, conversation, and commuting can also be cap-

tured. The list of behaviors that can be inferred from sensory data is growing rapidly.

Paraphrasing [9], we define **context** as any information that can be used to characterize the situation of a behavior. As with behaviors, contexts are explicitly stored in, or inferred from, measurements. The example contains four types of contexts which are representative of the capabilities of today’s personal sensing systems: temporal, physical, psychological, and social.

Temporal contexts describe characteristics related to the timing of a behavior, such as the exact start time of a behavioral episode. In the example, the temporal properties include the start and end timestamps of Jane’s run, seizure, emotional states, and time with her husband. **Physical** contexts describe the physical environment in which a behavior occurs, such as location and objects at a location. In the example, the addresses of the start and endpoint of the run, and the route taken between them, define these properties. **Psychological** contexts describe the psychological state of the user during the behavior. In the example, Jane experienced four psychological states: angry, stressed, and relaxed while running and terrified during the seizure. There is a wide gamut of emotions a person can experience, all of which fall under this category [20]. **Social** contexts describe the social environment in which a behavior occurs, and could include who the user was with when the behavior occurred and whether the user was interacting with that person. In the example, Jane’s social context was initially empty. Later, her husband John entered the social context. Note that contexts cannot only be associated with behaviors but also with other contexts. This is demonstrated in the association of timestamps (temporal), with Jane’s locations (physical), emotions (psychological), and the presence of her husband (social).

Privacy Threats

Privacy threats are the risks or harms that could come to the data producer if the his or her identity is associated with the data. If identity were removed from the dataset, the number of threats to the data producer decreases. Unfortunately, maintaining identity privacy is not always feasible with sensory datasets for two reasons. First, identity is sometimes needed to retain the utility of these datasets. A caregiver could develop a personalized treatment for a patient from his/her dataset, but would not know who to give the treatment to without the patient’s identity. Second, identity is often intertwined with useful information that the data producer would like to share. In Jane’s case, reidentification is possible from the GPS data [19] or from the presence of John’s identity in the dataset.

Once a dataset is re-identified, there are three types of threats that could emerge: financial, psychological, and physical. **Financial threats** are threats that lead to loss of assets or property. It includes professional threats, such as the loss of a job or damage to one’s business reputation. **Psychological threats** affect the data producer’s emotions. Such threats include embarrassment due to demasking of white lies or demasking of emotion regulation, deterioration in social

or family relationships, and development of pathological psychological conditions. **Physical threats** are threats to personal safety that may result in physical harm to the data producer. To reduce the probability of these threats when identity privacy cannot be maintained, we must maintain behavior and context privacy using restrictions and abstractions.

Restrictions and Abstractions

Restrictions remove data from a dataset before it is shared to reduce the potential privacy threats in the data. Measurements, behaviors, and contexts can all be restricted. For example, to prevent her husband’s identity from identifying her own, Jane could restrict access to her bluetooth data. Additionally, to make it difficult to infer the seizure, Jane could restrict access to accelerometer data. Note, however, that this would also prevent the AutoSense team from accessing Jane’s physical activity data, an important factor in moderating stress that may be of scientific interest.

Restricting all measurements, behaviors, and contexts would effectively empty the dataset, making it useless. On the other hand, sharing all the data is not desirable either. **Abstractions** provide a middle ground. They operate on measurements, behaviors, and contexts to reduce the extent of the exposure in the dataset. Abstraction operators include suppression[32], substitution[26], and transformation[29].

Jane’s privacy could be better protected using several abstractions. John’s presence during Jane’s run and seizure (social context) could be abstracted into “family member present.” Jane’s terrified state during the seizure (psychological context) could be abstracted into a less specific emotion such as “stressed.” This would not remove the seizure from the dataset, but would make it more difficult for an adversary to search for unusual emotional events like the seizure.

The examples above imply that one can mitigate the threats associated with a deanonymized dataset by restricting or abstracting behaviors and contexts. Stated mathematically, the probability of a set of threats \bar{T} is $P(\bar{T}) = f(\bar{R}, \bar{A}, \bar{M})$, where \bar{M} is the set of measurements, and \bar{R} and \bar{A} are sets of restrictions and abstractions, respectively, applied to \bar{M} or behaviors \bar{B} and contexts \bar{C} inferred from \bar{M} . Next, we describe a study that assesses how real people perceive $P(\bar{T})$ under various combinations of \bar{C} , \bar{B} , \bar{R} , and \bar{A} .

STUDY DESIGN

To study the privacy concerns associated with the disclosure of measurements collected by wearable sensors in the mobile environment, we designed a user study with three goals.

Goal 1: Assess the privacy concerns of real people regarding disclosure of continuously-collected physiological, behavioral, and psychological data. We also assess the change in concern levels as personal stake in the data is increased [24].

Goal 2: Use the proposed privacy framework and examine how restrictions and abstractions applied to various behaviors and contexts change concern levels.

Goal 3: Assess how reidentification of the data producer affects concern levels as the type of data consumer is varied.

Participants

66 participants were recruited from the student population at a 20,000+ student university in the United States. Participants were recruited using flyers and word-of-mouth. Participants volunteered to join one of two groups, a group with no personal stake in the data (Group *NS*) or a group with a personal stake (Group *S*) in the data. Table 1 summarizes the demographic characteristics of the groups.

	N	Males	Age±Std	AP	B	W	O
NS	36	56%	25 ± 4	22%	22%	53%	3%
S	30	53%	23 ± 4	23%	33%	40%	3%

Table 1. Group NS and S Demographics: AP=Asian/Pacific Islander, B=Black - Non Hispanic, W=White - Non Hispanic, O=Other

To capture the concerns of people with a personal stake in such data, we integrated this study into an existing study whose participants wore the AutoSense system for three consecutive days (Group *S*). The existing study examined the use of micro-incentives for scientific data collection and the effect of interruptions on user stress level. The data on micro-incentives and interruptions is outside the scope of this article and will be reported separately.

Procedure

For three days, Group *S* participants ($N = 30$) collected physiological, behavioral, and psychological data using the AutoSense sensor system as they went about their normal everyday life. At the end of the three-day period, Group *S* participants completed a privacy questionnaire assessing their concern regarding disclosure of selected behaviors and contexts with various restrictions and abstractions applied. Next, they reviewed graphs depicting the various behavioral and contextual inferences that were derived from their data (e.g., periods of stress, conversation, commuting, etc.) to further increase their personal exposure to and overall stake in their data. Finally, they completed the privacy questionnaire again. This within-subjects, repeated measures design allowed assessment of Group *S* concern level (Goals 1 and 2) at both a low (pre-review) and high (post-review) level of personal stake in the data (Goal 1). To distinguish between these two levels, the rest of this article refers to Group *S* before the review as Group *S-Pre* and Group *S* after the review as Group *S-Post*.

Group *NS* participants ($N = 36$) had no exposure to continuous physiological, behavioral, and psychological data collection. They did not wear AutoSense and did not review any data collected by it. They only completed the same privacy questionnaire as Group *S*. This allowed a between-subjects comparison of concern levels between participants with no personal stake in the data, Group *NS*, and participants with a personal stake in the data, Group *S* (Goal 1).

Group *NS*

Participants in Group *NS* first provided informed consent. After consent was given, they completed the demographics

questionnaire followed by the privacy questionnaire. They were then paid \$2 for participation.

Group *S*

Participants in Group *S* first provided informed consent. As part of the informed consent process, Group *S* was informed about the data collected by the AutoSense sensors. After consent was given, they completed the demographics questionnaire. Then the field protocol was explained to the participant. After explaining the protocol, the study coordinator demonstrated how to put on the sensors and then assisted participants in putting them on if necessary. The study coordinator verified the sensors were working properly by visually examining the streaming sensor data using an oscilloscope program. Once the sensors were verified as working, the participant was sent into the field.

Participants wore the AutoSense sensors for 3 consecutive days during their awake hours. They were instructed to take the system off at night and put it back on in the morning. Participants also carried the AutoSense mobile phone (Android G1). Periodically, the phone would ask participants to complete questionnaires. Participants earned micro-incentives (between \$0.02 and \$0.11) for each question they completed.

On the last morning of the study, the participant returned to the lab where he/she completed the privacy questionnaire. After the questionnaire, the participant reviewed graphs depicting the data he/she collected over the 3 days (described below), and then completed the questionnaire again. Lastly, participants were debriefed to collect subjective comments about their data and any concerns they had about it.

Data Review Session

We developed the *Aha* visualization system for reviewing data collected by the AutoSense system from natural environments. *Aha* is designed to help participants examine their daily behaviors and contexts at low, medium, and high levels of abstraction. *Aha* incorporates four visualizations. The **Day at a Glance** visualization depicts an overview of behaviors performed by individuals in their daily life. For example, the visualization presents the fraction of time (temporal context, high abstraction) spent commuting or in conversation. Psychological context is also depicted here as durations of stress in a day (Figure 3). The **Stress at a Glance** visualization depicts the fraction of time (temporal, high) participants are stressed (psychological, medium) during behaviors such as commuting or walking. The **Stress at Places** visualization depicts the fraction of time (temporal, high) at places, such as home and work (physical, high), when participants are stressed (psychological, medium). Lastly, the **Daily Timeline** (Figure 4) visualization plots behaviors on a detailed timeline (temporal, low). An analysis of 100 hours of AutoSense data found that the visualizations in *Aha* were approximately 80.16% accurate. Group *S* participants reviewed the data they collected using *Aha* for 10-15 minutes.

Privacy Questionnaire

The privacy questionnaire asks participants to rate their level of concern regarding disclosure of various combinations of

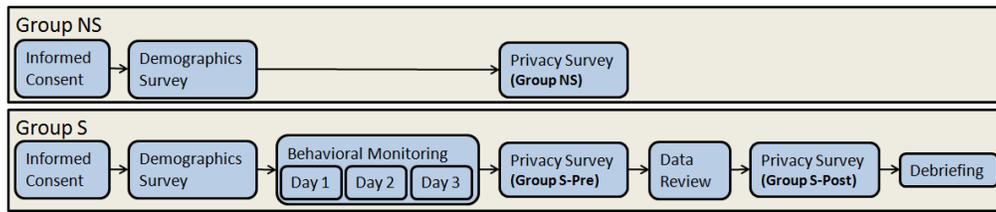


Figure 2. Procedure for Groups *NS* (top) and *S* (bottom)

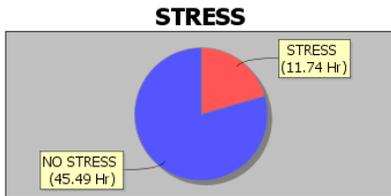


Figure 3. Fraction of monitoring period participant was stressed

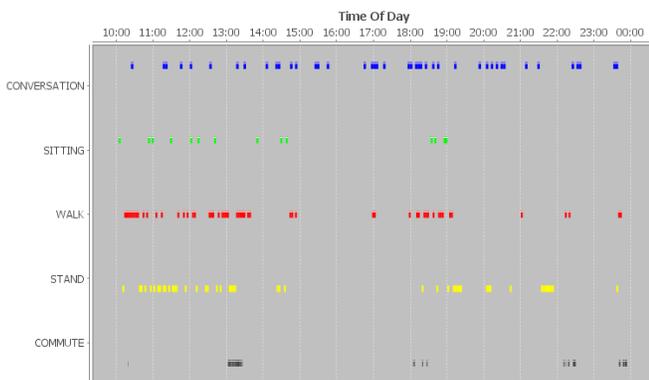


Figure 4. Daily Timeline of a participant monitored by AutoSense

behaviors, contexts, and levels of abstraction¹ (Goal 2). Concern is rated on a five-point scale: Not concerned (0), A little concerned (1), Moderately concerned (2), Concerned (3), and Extremely Concerned (4). The questionnaire consists of 9 sections. Sections 1 through 6 ask participants to rate their concern level for disclosure of places, smoking, stress, conversations, commuting, and exercise habits, all of which were depicted in the *Aha* visualizations. However, all participants reported they did not smoke, so the smoking section was automatically skipped by the questionnaire software.

Within each section, the participant is asked to rate the section's behavior or context if it were released with no context, temporal context, physical context, and both temporal and physical context simultaneously. The abstractions of temporal context on the questionnaire are timestamp, duration, and frequency. Physical context is assessed at one level of abstraction, place. Questions that ask about both temporal and physical context simultaneously ask about place and timestamp and place and duration. With each question, an example scenario of the disclosure is provided to help the

¹Available for download at [10]

participant understand the question. The last question for all behaviors, stress, and place asked participants to write in an explanation for ratings of at least moderately concerned. The choice of behaviors, contexts, and abstractions in the questionnaire was driven by 1) the capabilities of AutoSense and *Aha*, and 2) limiting the number of questions on the survey to approximately 50. The latter requirement was particularly important given the already high burden placed on Group *S* participants, who filled out the privacy questionnaire after 3 days of wearing sensors and answering field questionnaires.

Section 7 asks participants what information about their daily life they are particularly concerned about sharing. Sections 8 and 9 assess how concerned the participant would be if the data were shared, with or without identity, with the study coordinators, other study participants, other scientists and researchers, and the general public (Goal 3).

RESULTS

We analyzed participant data with respect to the three goals of the study as discussed in the preceding section. We now discuss the results and their interpretation from comments provided by the participants during debriefing. Unless otherwise noted, two-tailed t-tests were used to test for significant differences ($p < 0.05$). In comparisons between Groups *NS* and *S-Pre* (different populations), unequal variances were assumed. In comparisons between Groups *S-Pre* and *S-Post* (same population), a paired t-test was used. Shapiro-Wilk tests confirmed the normality of the data. Where space allows, p-values are reported with the means and standard deviations of the corresponding distributions.

Goal 1: Concern Levels and Personal Stake

To analyze participant concerns and the effect of personal stake on those concerns, per-participant averages of concern were calculated for conversation, commuting, exercise, stress, and places for all three groups. Figure 5 depicts these summary measures for all three groups.

Group *NS* and Group *S-Pre* provided similar concern ratings for exercise and place, and no significant differences were found between their distributions. However, there was a trend toward significance for commuting ($p = 0.08$, $NS = 1.3 \pm 0.1$, $S-Pre = 1.9 \pm 0.3$) and conversation ($p = 0.09$, $NS = 1.2 \pm 0.1$, $S-Pre = 1.4 \pm 0.2$).

Group *S-Post* had higher concern ratings than Group *S-Pre*. In particular, conversation for *S-Post* was significantly higher ($p = 0.01$, $S-Pre = 1.4 \pm 0.2$, $S-Post =$

2.1 ± 0.2), and similar trends toward significance were found for exercise ($p = 0.0978$, $S-Pre = 0.968 \pm 0.1$, $S-Post = 1.0 \pm 0.5$), commuting ($p = 0.07$, $S-Pre = 1.9 \pm 0.3$, $S-Post = 2.4 \pm 0.3$), and stress ($p = 0.06$, $S-Pre = 1.8 \pm 0.4$, $S-Post = 2.6 \pm 0.3$). Furthermore, across Group $S-Post$, 20% of stress questions were rated at “Concerned”(3) or “Extremely Concerned”(4), 15% for commuting, 12% for conversation, and 8% for exercise. Disclosure of place did not see a significant increase in concern from pre- to post-review (similar to that observed in [4] upon suitable scaling).

Taken together, the results for Groups NS , $S-Pre$, and $S-Post$ indicate increasing personal stake in the data helps people better estimate their concerns regarding the disclosure of behaviors and contexts. The highest level of concerns emerged for exercise, conversation, commuting, and stress after participants observed visual depictions of their data. Surprisingly, only 10-15 minutes of data observation was needed for these differences to emerge.

To better understand the rationale for participants’ concern ratings, we examined their free-form survey responses and debriefing comments. Two participants indicated their concern regarding commuting stemmed from an overall concern about being watched: “I am uncomfortable with someone watching over me that closely.” Two other participants expressed similar concerns: “It would be strange knowing that a person knew exactly where I was going to be, and when, at all times of the day. I wouldn’t feel like I had any privacy.”

Although the survey did not have any questions about social context, some participants said they were concerned about revealing social context with conversation: “I don’t like people to know whom I talked to at all...as i feel it’s none of their business.” Another participant said, “I do not want to share even when I talk because people may find out with whom if it is marked with time, which can be private.” This participant felt that adding time to conversation data may provide enough information to reveal who he talks to.

Participant concerns about stress appear to stem from a fear of having one’s inner thoughts revealed. One participant said, “I have this level of concern because whenever I am stressed out, I am around my mother.” The comment implies the participant does not want to reveal to her mother that she is a source of stress. Another participant said: ‘I feel my stress data if revealed can be a professional issue as am generally hot headed.’ People expect psychological states and inner thoughts to be private. Psychological monitoring technologies threaten this expectation.

Although concerns about exercise were low, participants commented that they did not want their exercise habits revealed. One participant explained: “I guess I’m a bit self-conscious over how little I exercise; I’d hate to have that broadcast.”

Goal 2: Impact of Applying Restrictions & Abstractions

Figure 6 breaks down the effect of selected restrictions and abstractions on individual behaviors and contexts. Table 2 highlights the effect of restricting or abstracting temporal

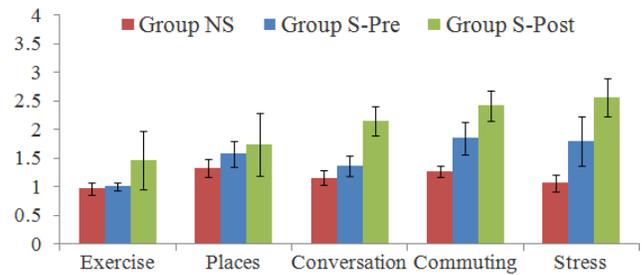


Figure 5. Groups NS and S concerns with respect to selected behaviors and contexts. The height of the bars represents the mean concern level, and the error bars depict \pm one standard deviation from the mean.

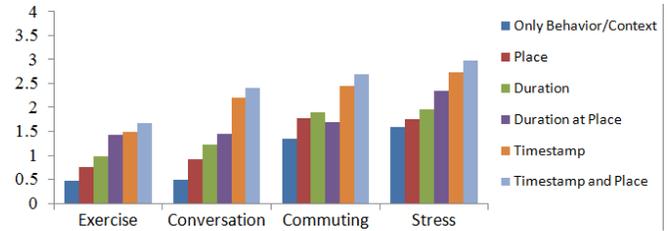


Figure 6. Changes in participant concerns due to temporal contexts

context. Generally speaking, adding temporal context to behaviors and other contexts increased concern level, with each decreasing level of temporal abstraction leading to a corresponding increase in level of concern. More often, adding timestamps induced a significant change in concern as compared to adding duration. Participants were most concerned about others knowing the exact moments they experience stress. Adding duration to stress events did not increase concern significantly, but adding timestamp increased concern by approximately 50%.

Participants were most concerned about sharing physical and temporal context together. There was a trend of increasing concern for place and timestamp, with the lowest concern for just sharing the behavior or context, the next level of concern for reporting the place or time of the behavior or context, and the highest level of concern for reporting both the place and time of the behavior or context. Across Group $S-Post$, 9% of behavior-only concerns were rated at “Concerned”(3) or “Extremely Concerned”(4), 11% for behavior with place, 20% for behavior with timestamp, and 26% for behavior with both timestamp and place. Adding both place and time simultaneously increased concerns significantly for conversation, commuting, and stress. Indeed, one participant explicitly expressed a concern about stress, time, and place: “I’d be concerned with people being able to tie the ‘whens’ of my being stressed with the ‘wheres.’ I’d rather people not know that I felt stressed at my particular job or when at my house, because they wouldn’t have the whole picture.” Exercise, however, did not exhibit a significant change from adding timestamp or place alone. Participants are likely less concerned about exercise episodes, since exercising is a positive behavior which they may be proud to share with others. Lower concerns for sharing place alone is consistent with the growing use and acceptability of location sharing sys-

	Exercise	Conversation	Commuting	Stress
Behavior/Context	0.02**, ns	ns, 0.06*	ns, 0.012**	ns, 0.06*
B/C + Place	0.008**, 0.08*	ns, 0.016**	ns, 0.062*	ns, 0.002**

Table 2. Effect of adding duration and timestamp to behaviors and other contexts (Group *S-Post*). Each cell is the result of a paired t-test comparing the particular row-column pair with that same pair, but with duration or timestamp added. The result for duration is the first value in each cell, and the second value is the result for timestamp. (ns = not significant, * = trend toward significance, ** = significant)

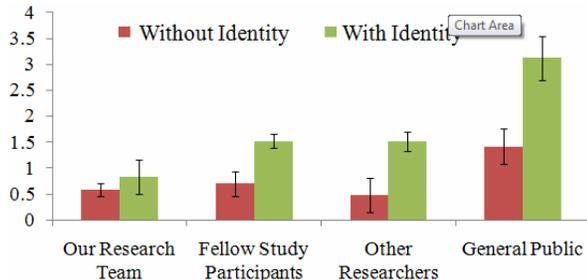


Figure 7. Change in participant concerns with respect to disclosure to different groups of people with and without identity (Group *S-Post*).

tems and what has been observed in other studies [4].

Goal 3: Role of Data Recipient

Participants rated their concern with respect to sharing identified and anonymous data with 4 groups: our research team, other study participants, other scientists and researchers, and the general public. Figure 7 summarizes participant responses. Adding identity to the dataset significantly increased participant concerns for all groups except our research team (Fellow Participants: $p=0.02$, $NoId = 0.7 \pm 0.2$, $Id = 1.5 \pm 0.1$, Other Researchers: $p=2E-4$, $NoId = 0.5 \pm 0.3$, $Id = 1.5 \pm 0.2$, General Public: $p=4E-5$, $NoId = 1.4 \pm 0.3$, $Id = 3.1 \pm 0.4$). 10% of participants rated releasing data to the general public without identity at “Concerned” or “Extremely Concerned.” This increased to 70% after adding identity. Releasing the data to members of our research team with identity increased concerns, but not significantly.

There was a significant increase in concern from releasing identified data to other researchers to releasing identified data to the general public ($p=0.001$, $OtherResearchers = 1.5 \pm 0.2$, $Public = 3.1 \pm 0.4$). Note that releasing identity to the general public more than doubles concerns to 3.1, the highest average concern level reported on the privacy survey. A trend toward significance was found between releasing identified data to members of our research team and releasing identified data to other study participants ($p=0.06$, $OurTeam = 0.8 \pm 0.3$, $OtherParticipants = 1.5 \pm 0.1$). Releasing identified data to study participants and to other researchers generated approximately the same level of concern. For anonymous data, increase in concern rating trended toward significance only when the recipient was changed from other researchers to the general public (0.06, $OtherResearchers = 0.5 \pm 0.3$, $Public = 1.4 \pm 0.3$).

The results indicate disclosure of any data to the general public is of significant concern to participants. When identity is added to the dataset, this concern more than doubles. Fur-

thermore, identity has a similar effect for release to fellow study participants or other researchers.

IMPLICATIONS AND LIMITATIONS

Awareness of Threats

The results highlight an overall lack of awareness of the actual information contained in personal sensing datasets and the privacy threats associated with them. Participants who had little or no personal stake in the data, representative of most of the population today, did not appear to understand the sensitive nature of the data.

Interestingly, concerns regarding place did not change after the review session. One interpretation is that participants were already aware of the potential threats associated with sharing place, due to the growing use of location-based services and media coverage of the issues associated with them.

Reducing Threats

The results imply the components of the proposed framework – behaviors, contexts, restrictions, and abstractions – can be manipulated to mitigate perceived threats in a dataset. Avoiding combinations of temporal and spatial context significantly reduced concerns about the dataset. Likewise, the results suggest one should be extremely careful with revealing psychological states, as indicated by the high concern regarding stress. Abstractions also clearly play a role in defining the threats in the data. Timestamps generally raised concerns significantly when combined with other contexts and behaviors. Increasing the level of abstraction (e.g., releasing duration instead) decreased these concerns.

Ultimately, the choice of which behaviors and contexts to release, and how they should be abstracted, cannot be decided purely based on the threats created by those choices. Increasing abstraction may reduce the threats in the data, but it also reduces the utility of the data. For example, a scientist cannot study daily variations in stress if only the durations of stress events are revealed. Likewise, a caregiver cannot study a patient’s gait if the patient’s accelerometer data is abstracted into “moving” and “not moving” states. Clearly, manipulations of restrictions and abstractions need to be tailored for both the data contributor and consumer. Only by customizing privacy transformations can the privacy requirements of the data producer and the information quality requirements of the data consumer be met. This same observation inspired Iachello and Abowd’s [14] complementary privacy framework, which aims to balance the utility of a UbiComp application with its burden on privacy. Personal sensing application developers can use our framework and study results within Iachello and Abowd’s to evaluate

the appropriateness and adequacy of various privacy transformations for the data contributor and consumer.

Threats from Revealing Psychological Context

The results of the study imply psychological context requires the most attention from the privacy and broader HCI community. Ratings indicated participants were more concerned about episodes of stress being revealed than any other behavior or context. The section on stress in the privacy questionnaire garnered more written comments than any other section, both pre- and post-review.

Psychological states like stress are different from behaviors and other contexts because they cannot be observed with the naked eye. Thus, psychological states are private by default, and remain private unless they are revealed. Personal sensing technologies that assess psychological states change this fundamental expectation. They enable mind-reading, a significant concern highlighted by participant comments.

However, such information can be useful, especially in the context of studying and treating psychological disorders (e.g., addiction and social anxiety). Methods to mitigate the harms associated with psychological context while still preserving their utility should be developed to address these concerns.

Limitations

A potential limitation of the study is that it confounds two variables that could affect concern level, 1) the personal nature of the data reviewed by Group *S*, and 2) increased awareness due to exposure to visual depictions of the data. Either variable alone or a mix of both could have increased concerns. A future study will address this, wherein Group *NS* will be shown visual depictions similar to those shown to Group *S*, except the depictions will not represent the participant's personal data. Assessing privacy concerns before and after exposure to the non-personal data will allow separating the effect of both variables.

CONCLUSIONS AND FUTURE WORK

Data collected by wearable sensors provide personal and societal utility, but also contain many unknown threats. Our study provides three key insights into the perceptions of threats from such datasets and how those threats can be mitigated. First, our results indicate people cannot understand the potential threats in the data unless they have a personal stake in it. Of the behaviors and contexts we examined, people were most concerned about revealing conversation, commuting, and inherently private psychological states (in this study, stress). Second, adding physical and temporal context increases concerns about the data, but those increases can be mitigated through restriction and abstraction. Third, people are willing to share such data - even with their identity - with us, other study participants, and other researchers. However, participants have significantly more concerns regarding sharing with the general public (e.g., over the web), especially when released with identity.

Given the concerns about psychological context, and our growing ability to capture psychological states, the com-

munity should examine disclosure of psychological context more closely. This paper did not assess participant concerns for combinations of psychological context with other behaviors, nor did it assess multiple abstractions of psychological context. Furthermore, it only assessed concerns for stress, which is not as specific as releasing specific emotions. In addition, the community could examine social context's effect on a dataset. Threats from social context are different than the others assessed here because they affect not only the user wearing the sensors, but also potentially any person interacting with the user.

Lastly, this study only examined privacy concerns from the perspective of data producers. The community should also examine how data consumers perceive privacy issues and what aspects of the data make it useful. Such a study would provide a better understanding of how to tradeoff behavior privacy and utility for physiological, psychological, and behavioral data collected by personal sensors.

ACKNOWLEDGEMENTS

We acknowledge the contributions of the following individuals. Anind Dey, Mustafa al'Absi, and Deepak Ganesan contributed to the study design, while Mary Read, Kelly Peck, Andrew Hoff, Ken Ward, and Satish Kedia helped conduct the study. Emre Ertin, Nathan Stohs, and Siddharth Shah developed the AutoSense sensor suite. Dan Siewiorek, Asim Smailagic, Patrick Blitz, Brian French, and Scott Fisk contributed to the study smartphone software. Context detection algorithms used in the visualization software were developed by Somnath Mitra (significant places from GPS data), Md. Mahbubur Rahman and Amin Ahsan Ali (speaking from respiration data), and Kurt Plarre (sensor detachments). This work was supported in part by NSF grants CNS-0910878, CNS-0910706 funded under the American Recovery and Reinvestment Act of 2009 (Public Law 111-5) and NIH Grant U01DA023812 from NIDA.

REFERENCES

1. AliveTec. Activity and Heart Monitor. <http://www.alivetec.com/products.htm>.
2. AutoSense. AutoSense: A Wireless Sensor System to Quantify Personal Exposures to Psychosocial Stress and Addictive Substances in Natural Environments. <http://sites.google.com/site/autosenseproject>.
3. M. Barbaro, T. Zeller, and S. Hansell. A Face is Exposed for AOL Searcher No. 4417749. *New York Times*, 2006.
4. L. Barkhuus and A. Dey. Location-Based Services for Mobile Telephony: a Study of Users' Privacy Concerns. In *Proc. Interact*, 2003.
5. A. Brush, J. Krumm, and J. Scott. Exploring end user preferences for location obfuscation, location-based services, and the value of location. In *ACM UbiComp*, pages 95–104, 2010.
6. S. Consolvo, J. Jung, B. Greenstein, P. Powledge, G. Maganis, and D. Avrahami. The Wi-Fi privacy

- ticker: improving awareness & control of personal information exposure on Wi-Fi. In *ACM UbiComp*, pages 321–330, 2010.
7. S. Consolvo, D. McDonald, T. Toscos, M. Chen, J. Froehlich, B. Harrison, P. Klasnja, A. LaMarca, L. LeGrand, R. Libby, et al. Activity sensing in the wild: a field trial of ubifit garden. In *ACM SIGCHI*, 2008.
 8. S. Consolvo, I. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location Disclosure to Social Relations: Why, When, & What People Want to Share. In *ACM SIGCHI*, 2005.
 9. A. Dey and G. Abowd. Towards a Better Understanding of Context and Context-Awareness. In *ACM SIGCHI workshop*, 2000.
 10. FieldStream. FieldStream: Network Data Services for Exposure Biology Studies in Natural Environments. <http://www.fieldstream.org/>.
 11. R. Ganti, N. Pham, Y. Tsai, and T. Abdelzaher. PoolView: Stream Privacy for Grassroots Participatory Sensing. In *ACM SenSys*, 2008.
 12. S. Guha, K. Plarre, D. Lissner, S. Mitra, B. Krishna, P. Dutta, and S. Kumar. AutoWitness: Locating and Tracking Stolen Property while Tolerating GPS and Radio Outages. In *ACM SenSys*, 2010.
 13. S. Hansell. AOL Removes Search Data on Vast Group of Web Users. *New York Times*, 2006.
 14. G. Iachello and G. Abowd. Privacy and proportionality: adapting legal evaluation techniques to inform design in ubiquitous computing. In *ACM SIGCHI*, pages 91–100, 2005.
 15. G. Iachello, K. Truong, G. Abowd, G. Hayes, and M. Stevens. Prototyping and sampling experience to evaluate ubiquitous computing privacy in the real world. In *ACM SIGCHI*, pages 1009–1018, 2006.
 16. W. Karim. Privacy Implications of Personal Locators: Why You Should Think Twice before Voluntarily Availing Yourself to GPS Monitoring, The. *Wash. UJL & Pol’y*, 14, 2004.
 17. P. Klasnja, S. Consolvo, T. Choudhury, R. Beckwith, and J. Hightower. Exploring privacy concerns about personal sensing. *Pervasive Computing*, pages 176–183, 2009.
 18. P. Klasnja, S. Consolvo, J. Jung, B. Greenstein, L. LeGrand, P. Powledge, and D. Wetherall. When I am on Wi-Fi, I am fearless: privacy concerns & practices in everyday Wi-Fi use. In *ACM SIGCHI*, pages 1993–2002, 2009.
 19. J. Krumm. Inference Attacks on Location Tracks. *Pervasive Computing*, 2007.
 20. R. Lazarus. *Stress and emotion: A new synthesis*. Springer Publishing Company, 2006.
 21. N. Li, T. Li, and S. Venkatasubramanian. t-Closeness: Privacy Beyond k-Anonymity and l-Diversity. In *IEEE ICDE*, 2007.
 22. K. Lorincz, B. Chen, G. Challen, A. Chowdhury, S. Patel, P. Bonato, and M. Welsh. Mercury: A Wearable Sensor Network Platform for High-Fidelity Motion Analysis. In *ACM SenSys*, 2009.
 23. A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. l-diversity: Privacy beyond k-Anonymity. *ACM TKDD*, 2007.
 24. C. Mancini, K. Thomas, Y. Rogers, B. Price, L. Jedrzejczyk, A. Bandara, A. Joinson, and B. Nuseibeh. From spaces to places: emerging contexts in mobile privacy. In *ACM UbiComp*, 2009.
 25. D. McFarland. Respiratory Markers of Conversational Interaction. *Journal of Speech, Language, and Hearing Research*, 44(1), 2001.
 26. M. Mun, S. Reddy, K. Shilton, N. Yau, J. Burke, D. Estrin, M. Hansen, E. Howard, R. West, and P. Boda. Peir, The Personal Environmental Impact Report, as a Platform for Participatory Sensing Systems Research. In *ACM MobiSys*, 2009.
 27. A. Narayanan and V. Shmatikov. Robust De-Anonymization of Large Sparse Datasets. In *IEEE Symp. on Security and Privacy*, 2008.
 28. A. Narayanan and V. Shmatikov. De-Anonymizing Social Networks. In *IEEE Symp. on Security and Privacy*, 2009.
 29. S. Oliveira and O. Zaiane. Privacy Preserving Clustering by Data Transformation. *J. Information & Data Management*, 1(1), 2010.
 30. L. Palen and P. Dourish. Unpacking Privacy for a Networked World. In *ACM SIGCHI*, 2003.
 31. S. Patil and J. Lai. Who gets to know what when: configuring privacy permissions in an awareness application. In *ACM SIGCHI*, 2005.
 32. L. Sweeney. Achieving k-anonymity Privacy Protection Using Generalization and Suppression. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), 2002.
 33. E. Toch, J. Cranshaw, P. Drielsma, J. Tsai, P. Kelley, J. Springfield, L. Cranor, J. Hong, and N. Sadeh. Empirical Models of Privacy in Location Sharing. In *ACM UbiComp*, 2010.
 34. E. Toch, N. Sadeh, and J. Hong. Generating default privacy policies for online social networks. In *ACM SIGCHI extended abstracts*, pages 4243–4248, 2010.
 35. J. Tsai, P. Kelley, P. Drielsma, L. Cranor, J. Hong, and N. Sadeh. Who’s viewed you?: the impact of feedback in a mobile location-sharing application. In *ACM SIGCHI*, 2009.
 36. Zephyr Technology. BioHarness BT. <http://www.zephyr-technology.com/bioharness-bt.html>.