EAI Endorsed Transactions
on Context-aware Systems and Applications
Research Article ICST.ORG

To Sense or not to Sense: An Exploratory Study of Privacy, Trust and other related concerns in Personal Sensing Context-aware Applications

Preeti Bhargava[1,*], Nick Gramsky[1], Ashok Agrawala[1]

[1]Department of Computer Science, University of Maryland, College Park, MD, USA

Abstract

Due to increasing proliferation of smart devices, many users store a significant proportion of personal data on them. Thus, personal sensing applications that sense a user's context via his smart device have significant privacy implications. In this paper, we conduct an exploratory study of privacy, trust, risks and other concerns of users with smart phone based context-aware personal sensing systems and applications. Our study results show that users are concerned that their sensed data can be misused, used for personal identification and tracking or for commercial purposes. However, they are willing to trade privacy for additional benefits if their sensed information is used for effective and beneficial causes. Furthermore, they are willing to trust reputed technology companies, with their data, if the benefits are significant. Based on these results, we propose a few design guidelines for designers of personal sensing apps and outline some interesting directions for future research.

Received on 15 May 2016; accepted on 27 June 2016; published on 12 September 2016

Keywords: Smart phone sensing; Personal sensing; Context-awareness; Privacy; Trust; Brand Recognition; Brand Awareness; Design guidelines

doi:10.4108/eai.12-9-2016.151676

1. Introduction

With the advent and ubiquity of smart devices such as smartphones and tablets, that come equipped with an increasing range of sensory, computational, storage and communication capabilities, a number of applications (also referred to as 'apps') and systems that can **sense** the user have emerged. This includes 'Hard Sensing' carried out through hardware sensors as well as 'Soft Sensing' done via application access or content extraction. These applications often focus on different types of context and activity recognition such as indoor and outdoor detection, physical activity recognition, and localization [1–5] and in different application areas such as location based services, social networking, health care etc. Commercial examples of such apps include Google Now[1] and Tempo[2] that sense and model a user's behavior based on his browsing history, emails and calendar data in order to provide him with personalized and relevant content.

Moreover, due to the increasing proliferation of these devices, many users carry them around all the time and perform a majority of their day to day activities (such as web browsing, listening to music etc.) via them. In addition, users store a significant proportion of personal data such as photographs, text messages, emails, calendars, and financial information etc. on the phone. As a result, personal sensing applications that track a user's context (such as his location, activities, behavior, browsing history and calendar content) via his smart device have greater privacy implications than traditional personal computers based applications as well as sensing applications that run on closed, proprietary devices such as Fitbits.

In this paper, we conduct an exploratory study of privacy, trust, risks involved and various other related concerns of users with such context-aware personal sensing systems and applications. In particular, we

*Corresponding author. Email: prbharga@cs.umd.edu
[1]http://www.google.com/landing/now/
[2]http://tempo.ai/

study several behaviors of users including those pertaining to:

- Their general privacy concerns with personal sensing applications and apprehensions about misuse of their sensed data.
- Data sharing - Their willingness to share data with other users, friends on a social network and other software.
- Sensitive Data Collection - Their willingness to allow the sensing app or system to collect sensitive health data as well as their perceived trade offs involved in storing this data on their smartphones vs on a cloud or server.
- Benefits to users - Their willingness to use a sensing app or system that stored sensed data on a server or cloud if it made smart decisions for them which had strong quantized benefits in terms of saving them time and money.
- Brand Recognition and User awareness - Their usage of services such as email, navigation etc. provided by a major technology company and their awareness about the company tracking their location, emails, as well as search, browsing, and video history.
- Brand Trust - Their willingness to use services provided by the aforementioned major technological company, despite the knowledge that their information is tracked and stored, if the services had strong quantized benefits such as saving them time and money
- Brand Reputation - Their willingness to use a sensing app if it were developed by a major technology company instead of a research prototype and saved them time and money.

We report results obtained from a live deployment with a smart phone sensing application and a web-based study involving 70 participants in all. Our results show that users are concerned that their sensed data[3] can be misused, used for personal identification and tracking or for commercial purposes. They are also concerned that the system or app may have unauthorized access to sensitive content on their devices and may be sharing their data with an external third party or sending it to a cloud or server. Moreover, the users want more control of what data they want to share, where it should be stored and how it should be mined. However, they are willing to trade privacy for additional significant benefits or if their sensed information is used for effective and beneficial causes. In addition, they are willing to trust reputed technology companies, which have a brand name, with their data if the benefits are significant despite being aware that their data is sensed and collected by these companies. Based on these results, we propose a few design guidelines for designers of personal sensing apps and outline some interesting directions for future research. To the best of our knowledge, other papers have not addressed such a broad spectrum of concerns with personal sensing applications.

The rest of the paper is organized as follows: Section 2 describes the methodology used for conducting our evaluation. Section 3 describes results extracted from the evaluation and Section 4 explains design guidelines inferred from the results. Section 5 discusses limitations of our work. Finally, we discuss related work in Section 6 and conclude in Section 7.

2. Methodology

The evaluation and results presented in this paper come from two studies:

- SenseMe system user study - We conducted exit interviews with 15 subjects after two, 2-week long live deployments of the SenseMe system.
- Web-based personal sensing privacy study - We conducted web based surveys among a population of 55 subjects that used or were aware of several smart phone based personal sensing applications but did not use SenseMe or take part in its user study.

We briefly describe these two studies now.

2.1. SenseMe System User Study

SenseMe[1] is an Android based system that leverages the smartphone and its various sensors such as accelerometer, GPS, WiFi, and Bluetooth in order to perform continuous, on-device, and multi-dimensional context and activity recognition for a user. It achieves this in a robust, automated, accurate, scalable, power efficient and non-invasive manner. SenseMe captures the following dimensions of a user's situation:

1. *Environmental context* - whether the user is outdoors (outside a building), indoors (inside a building), or indoor-outdoor (inside a building - near the door or a window),
2. *Location* - indoor/outdoor locations and type of location,
3. *Physical Activity* such as Walking/Running etc,
4. *Device Activity* - the task the user is currently engaged in on his/her smart device (checking mail, phone call)
5. *Social context* - how many people are around the user.

In SenseMe, all computation and processing is carried out on the device without requiring an external server. Moreover, the users' data is kept private and confidential on their devices and is visible only to them in order to mitigate privacy concerns.

[3] We use information and data interchangeably in this paper to refer to the users' sensed information.

Category	Pertinent question	Type
App Installation	"Before installing any smart phone application, do you read the EULA and privacy rules?"	Likert scale
Data misuse	"Are you concerned that the data sensed and collected by a smart phone sensing app could be misused?"	Likert scale
Privacy concerns	"What privacy concerns would you have with a smart phone sensing app?"	Free text
Data control	"If a personal sensing app allowed you to limit the data collected, what would you limit and why?"	Free text
Data Sharing	"If a smart phone sensing app shared your sensed data (such as activity or location) with other users of the aop in order to alert them that you are nearby (say for finding friends), would you use it?"	Likert scale
Data Sharing	"If a smart phone sensing app shared your sensed data (such as activity or location) with your friends on a social network that you used often, would you use it?"	Likert scale
Data Sharing	"If a smart phone sensing app shared your sensed data (such as activity or location) with other software, services or systems for user modeling purposes, would you use it?'	Likert scale
Data Storage and Retention	"Suppose the data sensed and collected by SenseMe or a similar smart phone sensing app was stored in a server or cloud (in an encrypted but unanonymized format). Would you use the app if it made smart decisions for you?"	Likert scale
Data Storage and Retention	"Suppose the data sensed and collected by SenseMe or a similar smart phone sensing app was stored in a server or cloud (in an encrypted and anonymized format). Would you use the app if it made smart decisions for you but not as effective as when the data wasn't anonymized?"	Likert scale
Sensitive Data Collection	"If the smart phone sensing app or system was able to sense and collect sensitive health data, such as heart rate, blood pressure, etc. while keeping this data on the phone, would you use it?"	Likert scale
Sensitive Data Collection	"If the smart phone sensing app or system was able to sense and collect sensitive health data, such as heart rate, blood pressure, etc. while sending this data to a cloud or server, would you use it?"	Likert scale
Sensitive Data Collection	"If the smart phone sensing app or system was able to sense and collect sensitive health data, such as heart rate, sweat rate, blood pressure, etc. while sending this data to a cloud or server and using it ONLY for saving lives, would you use it?"	Likert scale
Benefits to users	"Would you be willing to use SenseMe or a similar app if it stored sensed data on a server/cloud and made smart decisions for you that saved you 10 minutes of your time?"	Likert scale
Benefits to users	"Would you be willing to use SenseMe or a similar app if it stored sensed data on a server/cloud and made smart decisions for you that saved you an hour of your time?"	Likert scale
Benefits to users	"Would you be willing to use SenseMe or a similar app if it stored sensed data on a server/cloud and made smart decisions for you that saved you 1 % of your salary?"	Likert scale
Benefits to users	"Would you be willing to use SenseMe or a similar app if it stored sensed data on a server/cloud and made smart decisions for you that saved you 10% of your salary?"	Likert scale
Brand Recognition	"Do you use any of the following services: email, navigation, Personal Digital Assistant (PDA), Location based services (LBS), cloud storage and search, provided by a major technology company?"	Yes/No
User Awareness	"Are you aware that the major technology company, which is mentioned above, tracks your location, emails, search history, browsing history, video history, and location searches?"	Yes/No
Brand Trust	"Now that you are aware that this company has the ability to track so much information about you, would you be willing to use the services provided by it if they could save you 1% of your salary?"	Likert scale
Brand Trust	"Now that you are aware that this company has the ability to track so much information about you, would you be willing to use the services provided by it if they could save you 10% of your salary?"	Likert scale
Brand Trust	"Now that you are aware that this company has the ability to track so much information about you, would you be willing to use the services provided by it if they could save you a significant fraction of your salary?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned above?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned above and saved you 1% of your income?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned above and saved you 10% of your income?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned and saved you a significant fraction of your income?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned above and saved you 10 minutes of your time?"	Likert scale
Brand Reputation	"Would you be willing to use SenseMe or a similar app if it were developed by the major technology company mentioned above and saved you an hour of your time?"	Likert scale

Table 1. Privacy and trust related questions from our study questionnaire

Two user studies, each lasting 2 weeks, were conducted for evaluating SenseMe. The studies involved 15 participants from the USA and their ages ranged from 21 to 40 (μ =27.5). 46.7% of the participants were female and 53.3% were male. The participants reported a variety of occupations including software engineers, health and wellness coordinators, educators, post doctoral associates etc. However, the majority of the subjects were students. Self-reported completed levels of education ranged from college to doctorates.

In both the studies, SenseMe was installed on the subjects' personal devices in order to capture the

context and activity information in a real life practical scenario, thus, making the evaluation more effective. All the subjects were asked to run SenseMe on their devices, in the background, for a period of 2 weeks while going about their daily life. They were also asked to keep a journal of their activities, locations, environments and number of people around them throughout the day. This allowed them to present an actual portrayal of their day for an effective evaluation of the application.

On conclusion of the user study periods, each subject was interviewed to discuss which dimensions they found useful and interesting as well as to evaluate their user experience. They were also asked in detail about their privacy concerns with such personal sensing applications, the data it could sense, where the data should be stored, how it could benefit them etc. Results that focused on the system performance, accuracy and resource utilization of SenseMe are presented in [1]. This paper presents results on the participants' responses to questions on privacy, trust and other related implications of smart phone based personal sensing applications such as SenseMe.

2.2. Web-based Personal Sensing Privacy Study

55 participants, who used or were aware of several smart phone based personal sensing apps, were recruited via social media, emails and word of mouth to participate in a survey of 41 questions. Their ages ranged from 21 to 50 years (μ = 31.5) and their demographic distribution was as follows: 60% from USA, 18.2% from United Kingdom and 21.8% from India. 40.1% of the participants were female and 59.9% were male. The participants reported many occupations including researchers, engineers, full time graduate students, consultants, entrepreneurs etc. Some of the participants were home makers. Self-reported completed levels of education ranged from some college to doctorates.

To maintain consistency, these participants were first given a short description of SenseMe. They were then given the same privacy related questionnaire as the subjects in the SenseMe User Study. The goal was to survey a large number of subjects, with no firsthand experience with the application, as part of the same study.

3. Results

Table 1 shows the various categories of open-ended and quantitative questions pertaining to privacy, trust and other related issues from our two studies. A majority of the quantitative questions had responses on a Likert Scale ranging from 1 (Highest or Most Likely) to 8 (Lowest or Least Likely) while others were either free text based or dichotomous (Yes/ No). We present results for each now.

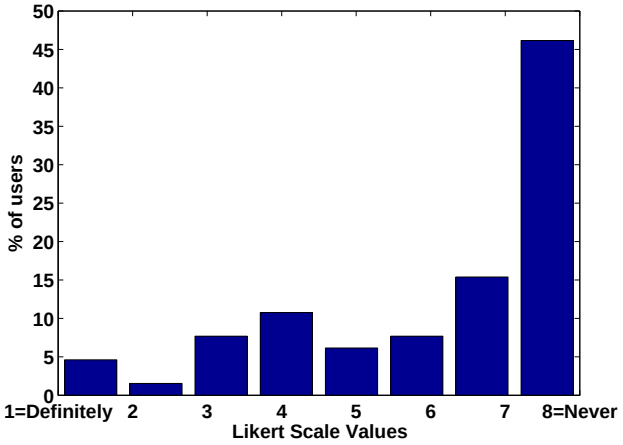

Figure 1. App Installation and EULA Responses

3.1. App Installation and EULA

We first investigated whether the subjects read the End Users' License Agreement (EULA) and kept track of what apps they were installing on their devices. Figure 1 shows the distribution (μ = 6.28, σ = 2.14) of responses of all the participants to the App installation question. More than 45% of the users never read the EULA (which includes the data, sensors and services on the phone that the app would access) while only 5% responded that they definitely read it. These results support the findings of Staiano et al. [6] that most users do not read the Terms of Service of smartphone apps. These findings also corroborate with those of Good et al. [7] with respect to users not reading computer software license agreements.

3.2. General Privacy Concerns

We first gauged whether the subjects had concerns about their sensed information being misused. The participants' responses had a μ = 1.1 and σ = 0.29. As evident by the low standard deviation and the high mean, all of the users were apprehensive that their sensed information could be misused and hence, expressed specific privacy concerns about its collection, monitoring and storage. Out of the 70 participants, 57 responded to this question. We applied the open coding method [8] to their responses, which is a standard method for analyzing qualitative data. We categorized the responses into several categories:

Concerns regarding sensed information being used for Personal Identification and Tracking. 18 of the 57 participants (31.6%) expressed concerns regarding the information being used for Personal Identification and Tracking. Some of the comments include:

- *"How much can be inferred about me from my app usage"*
- *"If people can hack it to tell when I am out of my house."*
- *"Whether or not 3rd parties could access data to determine patterns of life"*
- *"No Location Tracking"*
- *"If it is known that all members of a household are not at home by their locations, then it is possible that someone could use the system to find the best time to rob a house."*

Concerns regarding sensed information being used for Commercial Purposes. 6 of the 57 participants (10%) said that they did not want their sensed information to be used by companies for commercial purposes such as targeted advertising. One subject remarked *"I don't want to be shown ads based on what videos I watched"*.

Concerns regarding sensed information being used for Unintended or Undeclared Purposes. 4 of the 57 participants (7%) declared that they wouldn't use the app if it used their sensed information for purposes that weren't declared or intended by the app designer or provider.

Concerns regarding Unauthorized Access to sensitive information, phone sensors and services. 19 of the 57 participants (33.3%) expressed concerns regarding the app having unauthorized access to sensitive information such as phone book or photo gallery or to intrusive sensors such as GPS, camera and microphone. Some of the comments include:

- *"I also don't want it to access my contacts and call or message them."*
- *"Whether it has access and saves my telephone number, password of accounts directly synced on my mobile like gmail, bank accounts etc"*
- *"I don't want an app listening/seeing things around me."*

Concerns regarding Sharing or storing of sensed information with a Third party or on a cloud/ server. 16 of the 57 participants (28%) expressed concerns about their sensed information being shared with a third party, posted on a social network or stored in the cloud. Explicit comments include:

- *"The app should not send out any information like my location, my contacts, my passwords, etc. to any server."*
- *"Posting my data to Facebook or other social networking platform without my knowledge"*
- *"I'd want to know whether the data was transmitted to another machine for collection. I'd also like the ability to decline transmission of the data on a case-by-case basis (perhaps you're in a situation that you don't wish to be recorded)"*
- *"..if the app professed to save lives but also shared data collected in order to market ads to me, I wouldn't use it"*

Concerns regarding technical side effects. 1 of the 57 participants (1.7%) said that the app should not slow down the phone's performance or reduce the battery life.

No privacy concerns. 5 of the 57 participants (8.8%) said that they would have no privacy concerns and would use such apps only if they have very specific needs for them. One of these subjects stated that he shuts off all tracking sensors such as GPS and Wi-Fi as soon he leaves home.

3.3. Data control

As mentioned, most subjects expressed concerns regarding privacy and misuse of their sensed data. We then asked them if they would like to have more control of the data being sensed and if there is any data on their phone that they would limit and never allow an app to sense. 55 subjects responded to this question. 41 of the 55 participants (74.5%) wanted more control of their data and the ability to decide what should be monitored, where it should be stored, and how it should be used or mined. They also wanted the ability to delete the data when they wanted to, and limit or disable data sensing. 6 of the 55 participants (10.9%) said they would not limit the sensing and monitoring and instead focus on limiting what *they* stored or used on the device. They would also like to see *"stringent enforcement against abuse of information"*. As one participant mentioned that *"if a device is capable of recording the data, I assume it will"*. 8 of the 55 participants (14.5%) said that it would depend on many factors such as the data being sensed or collected and what it was being used for.

These 41 participants, who said they wanted to limit or disallow sensing, specified several types of information that they would not allow an app to sense. The specified information can be categorized into the following categories (the % indicate the fraction of the 41 users who specified items of this category):

- Personal identification data such as name (16%) or location (9%)
- Private or sensitive data such as photos and music (22%)
- Browsing and search history, chats (18%) and emails (16%)
- Calendar, notes and contacts (22%)
- Apps being used (4%)
- Calls (4%) or text messages content (18%)
- Access to sensitive sensors such as camera or microphone (5%) or services such as Wi-Fi or 3G plan (2%)

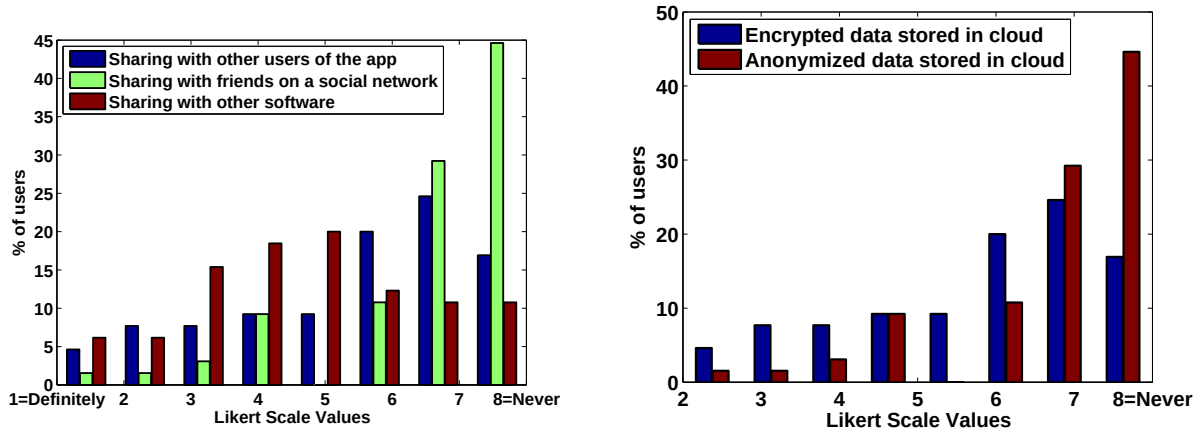

(a) Distribution of Data Sharing Responses

(b) Distribution of Data storage and Retention responses

Figure 2. Responses to Data sharing and Data Storage and Retention Questions (best viewed in color)

- Social networks data and video history (6%)
- Financial information such as bank accounts, SSN, stored passwords and online purchases (13%)

3.4. Data Sharing

We now investigate the subjects' willingness to share their data. Figure 2(a) shows the responses of the subjects when asked if they were willing to use an app if it shared their sensed data (such as activity or location):

- With other users of the same app - The responses had $\mu = 5.54$ and $\sigma = 2.1$.
- With their friends on a social network - The responses had $\mu = 6.92$ and $\sigma = 1.36$.
- With another software or service for user modeling purposes - The responses had $\mu = 4.74$ and $\sigma = 1.96$.

As evident from Figure 2(a), 45% of the subjects would never allow posting of sensed data such as activity or location on a social network, 15% responded that they would never share their sensed data with other users of the same app, and only 10% responded that they would never share their sensed data with another software for user modeling purposes. Thus, the subjects' willingness to share data with friends on a social network is lower than that for sharing data with other users of the same app or with another software.

3.5. Data Storage and Retention

Many of the subjects had expressed concerns regarding the sensed data being transported to another system or cloud over the network (see Section 3.2). However, for several context-aware and ubiquitous computing systems, a back end server side system is necessary. In such cases, the data transmitted over the network maybe either encrypted or anonymized but this can lead to a downgrade in performance.

Hence, we asked the subjects whether they would use such an app or system which transported and stored their data in an encrypted format in the cloud and made smart decisions for them. We also asked them if they would use the system if it anonymized the data but wasn't as effective as the system that only encrypted it.

Figure 2(b) shows the distribution of responses which are mildly positive. For a system that stored data in a cloud in an encrypted format but made smart decisions for the user, the subjects' responses had $\mu = 4.54$ and $\sigma = 1.96$. For a system that anonymized the data but wasn't as smart, the responses had $\mu = 4.55$ and $\sigma = 2.02$. Thus, the subjects were slightly more willing to risk the data being unanonymized than unencrypted, if it improved the system's performance in making smart decisions for them, though the difference is negligible.

3.6. Sensitive Data Collection

We next investigated the users' willingness to use apps or systems that collected and analyzed sensitive data but with certain tradeoffs and benefits. For this purpose, we asked them to rate their willingness to use an app that sensed their health information such as heart rate, blood pressure, etc. (which is highly sensitive information) and stored it on their phone, in the cloud and in the cloud but only used it for the purpose of saving lives.

Figure 3(a) shows the distribution of responses. If the data was stored on the:

- Phone only - μ was 3.42 and σ was 2.23.
- Cloud - μ was 4.74 and σ was 2.13,
- Cloud but used for the purpose of saving lives - μ was 3.52 and σ was 2.0.

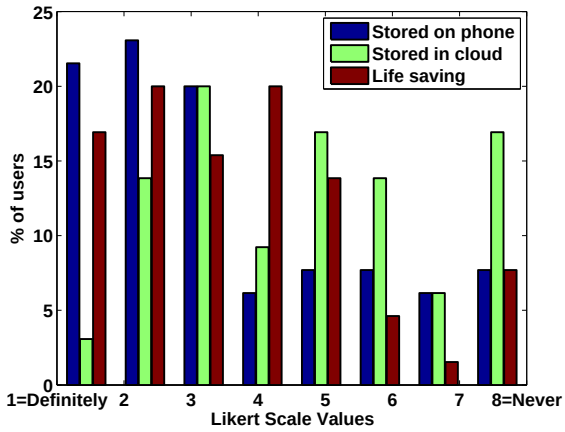

(a) Distribution of Responses to Sensitive Data Collection

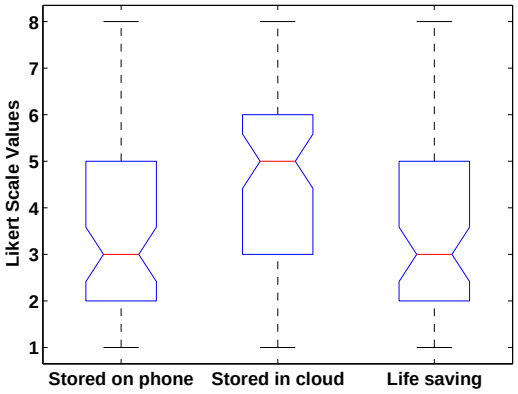

(b) One way ANOVA of responses to Sensitive Data collection

Figure 3. Responses to Sensitive Data Collection Question (best viewed in color)

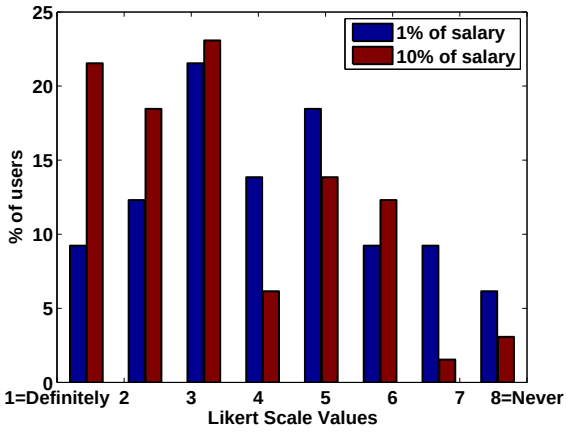

(a) Saving money

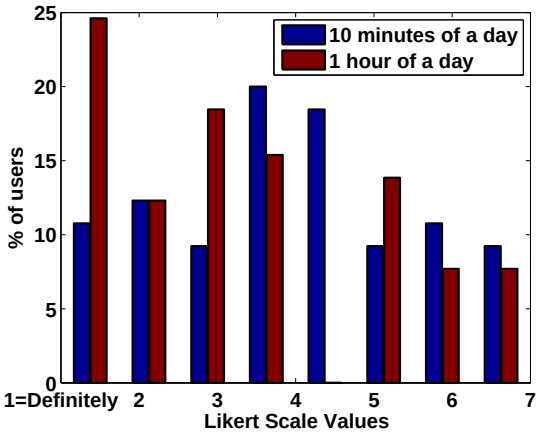

(b) Saving time

Figure 4. Responses to Benefits to users questions (best viewed in color)

In addition, we also performed a one-way ANOVA (with significance level α as 0.05) on the subjects' responses. Figure 3(b) shows the box plots for the subjects' responses. As evident from this distribution, the participants' willingness was higher when the data was stored on their phone. It decreased when the data was stored in the cloud but increased once again when it was mentioned that it will be stored in the cloud but used for saving lives. The mean values of the two box plots - for responses to health data being stored on the phone and responses to health data being on the cloud if it saved lives, are aligned very closely. This indicates that these distributions are not statistically different. Thus, if there is a tradeoff for the sensitive data to be utilized in a way that can prove beneficial, the users' willingness is higher and similar as compared to when the data is stored on their devices.

3.7. Benefits to users: Time=Money

As mentioned earlier, many subjects had concerns about their data being stored in the cloud or on a server but several context-aware and mobile systems require back end processing in order to be effective and efficient. We had investigated the subjects' willingness to use an app that stored their data in a cloud and made smart decisions for them in Section 3.5. However, in that case, the benefits were hypothetical. Here, we investigated their willingness to use SenseMe or a similar app if it stored data on a server or cloud and made smart decisions for users that had quantized benefits in terms of saving them time and money.

Saving money. Figure 4(a) shows the distribution of responses. When asked if the system saved them 1% of their salary, the responses had $\mu = 4.15$ and $\sigma = 1.99$.

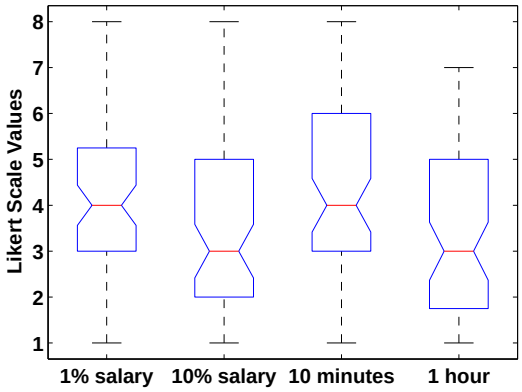

Figure 5. One way ANOVA of responses to User Benefits Questions

Intuitively, the willingness was higher ($\mu = 3.31$, $\sigma = 1.94$) when the amount of money saved was increased to 10% of their salary.

Saving time. Figure 4(b) shows the distribution of responses. If the system saved them 10 minutes of their time, the response had $\mu = 4.4$ and $\sigma = 2.11$. If the system saved an hour of their time, the willingness was higher with $\mu = 3.35$ and $\sigma = 1.92$.

Thus, as the hypothetical benefits increased in terms of time or money, user acceptance of storing sensed data on the cloud increased. In addition, we also we also performed a one-way ANOVA (with α level as 0.05) on the subjects' responses. Figure 5 shows the box plots for the subjects' responses. As shown, the mean values of the box plots for responses to saving 10 minutes and 1% salary are aligned very closely. Similarly, the mean values of the box plots for responses to saving 10% salary and 1 hour of time are aligned very closely. This indicates that these two response distributions are not statistically different. Thus, the subjects seem to view the benefits for time and money as equivalent and the greater the benefit, the higher their willingness to use the app or system.

3.8. Brand Recognition and User Awareness

So far in the study, we had been referring to SenseMe or a hypothetical smart phone app or system. However, several technology companies are building such systems already or have deployed similar systems in the real world where they continuously monitor and sense their users. Therefore, we wanted to evaluate whether the subjects recognized this, were aware of the technologies and brands, and if they might vary their behavior based on their awareness.

To this end, we investigated whether our study subjects recognized the brand of a major technology company and used various services (Mail, Navigation, Personal Digital Assistant (PDA), Location Based Services (LBS), Cloud storage and Web Search) provided by it. We also investigated if they were aware that their video history, location trace and location searches, and browsing and search history, were tracked by this company. In addition, we asked the subjects if they were aware that the company stores their data on servers and in the cloud.

Figure 6(a) shows the distribution of responses to the dichotomous questions related to usage of technologies and services provided by a major technology company. On an average, 61.54% of the subjects used at least one of the technologies or services. As shown, majority of the users were aware of and used the more popular services such as search, navigation, cloud storage and mail. The significantly lesser usage of the other two services - Personal Digital Assistant and Location Based Services could be because the former is available only on select devices while the latter is not so well known. Figure 6(b) shows the distribution of subjects' responses to being aware that the company tracked their video history, location and location searches etc. On an average, 88.97% of the subjects were aware of the company tracking their personal data.

This awareness of how major brands track personal data and the widespread acceptance of services that use this data seem to contradict earlier input from the subjects. All but one subject used the email service provided by the major tech company that we asked about in the survey. This email service tracks the social network that the emails themselves create, mines specific content in the email messages and uses it to for targeted advertising. Yet when we asked the subjects to freely list items a smart phone app should never sense, over 25% of respondents mentioned chats, text messages and emails. Additionally, many felt that using personal data for monetization purposes was wrong. This appears to show that users lower their convictions about private data with brands that they trust[4]. We will investigate this further in the future.

3.9. Brand Trust

Responses of all subjects. We then investigated whether the subjects would continue to use the services provided by the company despite being aware that their data was tracked along several dimensions. As an incentive, we again applied the hypothetical benefit of saving money, in order to see the trade-off between benefit and data sharing.

Figure 7 shows the distribution of responses. If the continued use saved the subjects:

[4]Although this compares an email service with a strict smart phone sensing app, email can also be a form of soft sensing.

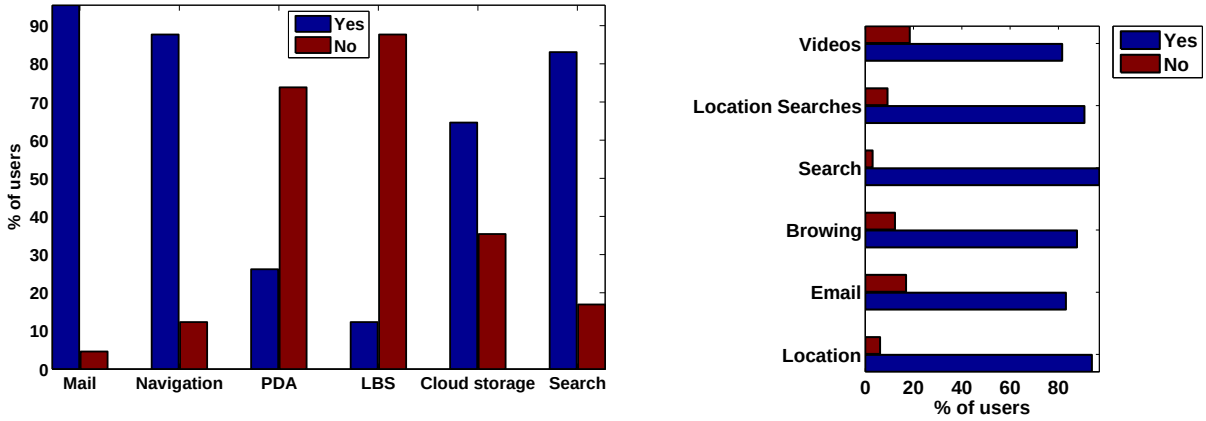

(a) Usage of technologies and services provided by a major tech company

(b) Awareness about tracking and sensing by a major tech company through its various services

Figure 6. Responses to Brand Recognition and User Awareness Questions (best viewed in color)

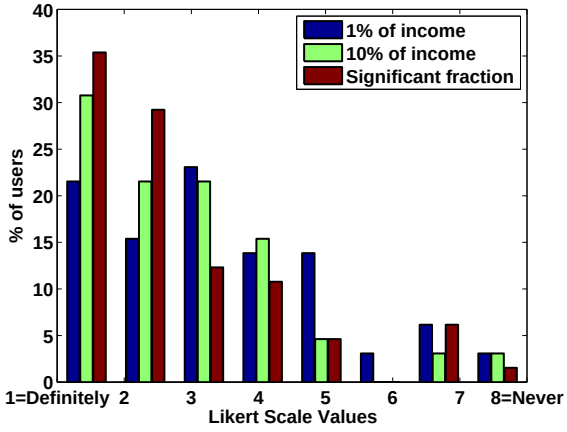

Figure 7. Usage of technologies and services provided by a major tech company despite being aware that users' data is sensed (best viewed in color)

- 1% of their salary - The distribution of responses had $\mu = 3.32$ and $\sigma = 1.92$.
- 10% of their salary - The distribution of responses had $\mu = 2.69$ and $\sigma = 1.73$.
- A significant fraction of the salary - The distribution of responses had $\mu = 2.52$ and $\sigma = 1.79$

Thus, the subjects were willing to use the technologies and services if the benefits were significant, despite being aware that their data was monitored. The % increase between the subjects' willingness to use these services if it saved them 10% of salary and if it saved them a significant fraction of it, is not very high. We believe that a possible reason for this could be that people are inclined to think that there is no concept of a 'free lunch'. Thus, if a service claims to save them a significant amount of money, they would actually be skeptical of it or mistrust it and dismiss it as fraudulent.

We next attempted to draw comparisons between these responses and the subjects' responses to the dichotomous Brand Recognition questions from Section 3.8 (regarding the usage of these technologies before being made aware that their data was tracked). To this end, we standardized the dichotomous results and Likert scale values to z-scores. Borenstein et al. [9] explain that it is possible to combine effect sizes from studies that used different metrics if there are comparable in relevant ways. Both the Brand Trust and Brand Recognition questions are same but the responses are on different scales. For the Brand Recognition question responses, we first converted the Yes and No responses for the 6 services to binary values, aggregated these binary values to convert them to a 6 point scale, and then converted the aggregated values to z-scores. For the Brand Trust questions, we directly converted the Likert scale response values to z-scores. The z-score is computed as: $z = \frac{x-\mu}{\sigma}$ where x is the raw value, μ is the population mean and σ is the population standard deviation. The distribution of the standardized responses for the usage of technologies:

- Before being made aware, had $\mu = 0.73$ and $\sigma = 0.67$.
- Being aware of the tracking and if it saved subjects' 1% of their salary, had $\mu = 0.81$ and $\sigma = 0.57$.
- Being aware of the tracking and if it saved subjects' 10% of their salary, had $\mu = 0.77$ and $\sigma = 0.63$.
- Being aware of the tracking and if it saved subjects' a significant fraction of their salary, had $\mu = 0.77$ and $\sigma = 0.63$.

Saving	Group A	Group U
1 % salary	$\mu = 3.28$, $\sigma = 1.95$	$\mu = 3.55$, $\sigma = 1.86$
10 % salary	$\mu = 2.67$, $\sigma = 1.77$	$\mu = 2.82$, $\sigma = 1.6$
Significant fraction	$\mu = 2.36$, $\sigma = 1.88$	$\mu = 2.56$, $\sigma = 1.29$

Table 2. Relationship between participants' awareness status and responses to the Brand Trust questions

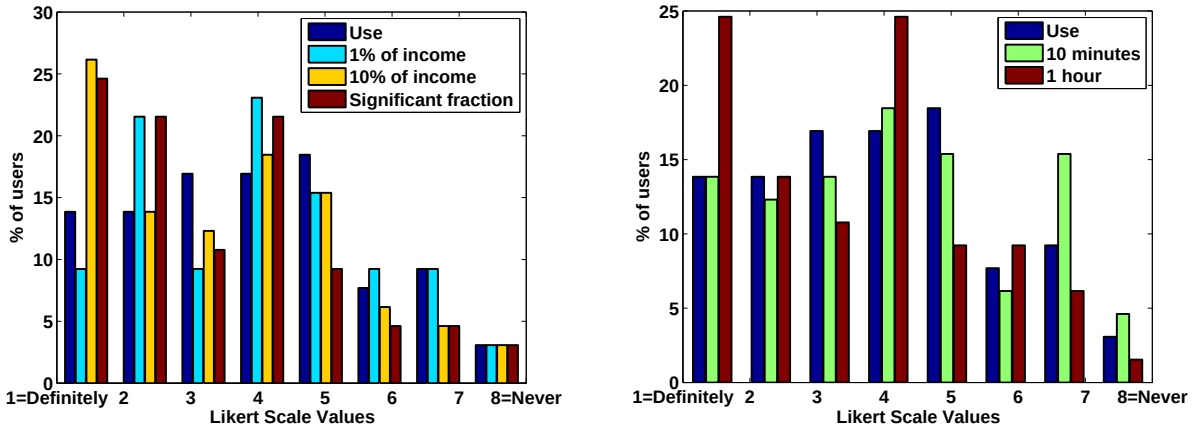

(a) Usage of an app similar to SenseMe if it were developed by a major tech company and saved them money

(b) Usage of an app similar to SenseMe if it were developed by a major tech company and saved them time

Figure 8. Responses to Brand Reputation Question (best viewed in color)

Thus, when compared to their responses earlier, the willingness is lower but increases slightly as the benefit increases.

Relationship between participants' awareness status and responses. To further investigate whether the new awareness impacted their decision to continue the use of technologies provided by this company, we divided the participants into two groups. These groups represented the participants' awareness status and were based on the z scores computed from their dichotomous responses to the User Awareness questions. The groups were the A group which consisted of subjects who had been aware of the tracking before we informed them (z score > 0.0) and the U group which consisted of the subjects who had been unaware of the tracking before we informed them (z score <= 0.0). We then compared the Likert scale responses, of these groups, to the Brand Trust questions. Table 2 shows the distribution. Clearly, the U group is more reluctant than the A group to use these technologies after being made aware.

3.10. Brand Reputation

We next analyzed whether the study subjects were willing to use an app similar to SenseMe if it were developed by the same major technology company and furthermore, had the same hypothetical benefits such as saving them time and money. Our intent was to determine if brand trust and reputation had a significant impact on their responses. Thus, our hypothesis was that since they already trusted the brand, their willingness to use the application and let it sense their data would be higher. As one subject from the SenseMe user study had mentioned in his exit interview, *"I would probably not use an app that monitors me through my smartphone, unless it's from a trusted source like Google."*

Responses of all subjects. Figure 8 shows the distribution of responses. We further compared the subjects' responses to this question with their responses earlier to the Benefits to Users (Section 3.7) questions where they were asked if they would use SenseMe or a similar app and share their sensed data if it saved them money and time. Figure 9 shows the results of one way ANOVA (with $\alpha = 0.05$) on the subjects' responses with respect to both money and time. For:

- Saving Money - If SenseMe saved them 1% salary, the responses had $\mu = 4.15$ and $\sigma = 1.98$ while for a similar product by a major technological company, the responses had $\mu = 3.94$ and $\sigma = 1.92$. Similarly, if SenseMe saved them 10% salary, the responses had $\mu = 3.31$ and $\sigma = 1.94$ while for a similar product, the responses had $\mu = 3.35$ and $\sigma = 1.99$.
- Saving Time - If SenseMe saved them 10 minutes in a day, the responses had $\mu = 4.4$ and σ

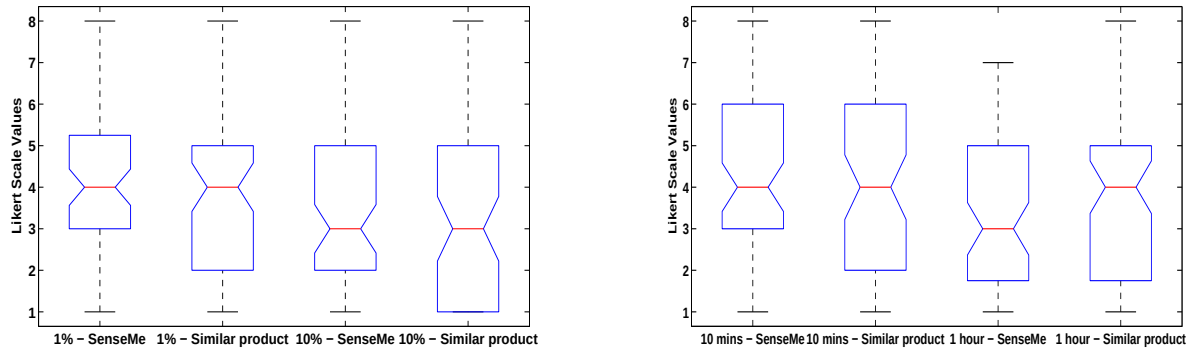

(a) Usage of SenseMe and usage of a similar app developed by a major tech company and both saved the subjects' 1% and 10% of salary

(b) Usage of SenseMe and usage of a similar app developed by a major tech company and both saved the subjects' 10 minutes and an hour

Figure 9. Anova for usage of SenseMe and a similar app developed by a major tech company and both saved the subjects' different amounts of money and time

= 2.1 while for a similar product by a major technological company, the responses had μ = 4.12 and σ = 2.1. Similarly, if SenseMe saved them 1 hour of a day, the responses had μ = 3.35 and σ = 1.92 while for a similar product, the responses had μ = 3.4 and σ = 1.96.

Thus, in both the cases, the willingness of all the subjects to use SenseMe was slightly lower than that for a similar app developed by the company mentioned. This seems to indicate that since the participants are not familiar with the reputation of the entity that developed the SenseMe app, i.e. our research group, they are less willing to trust us. On the other hand, the technology company mentioned is very well known and has a major brand so they are more willing to trust it.

Relationship between participants' group and responses. As mentioned in Section 2, the results come from two sources: the exit interviews of 15 subjects who participated in a live deployment of SenseMe, and web-based surveys filled by 55 participants. We next analyzed the differences in responses of these two groups - the L group which consists of subjects who participated in the live deployment and the O group which consists of the subjects who participated in the online surveys. Our intent was to reveal any influence that the first-hand usage of the app had on their responses.

A chi-squared test revealed that there was significant differences in the responses of the two groups. Table 3 shows the distribution. As shown by the μ and σ values for the Likert scale responses, the L group was more willing to use SenseMe than a similar app developed by a major technological company. On the other hand, the O group was more inclined to use a similar app developed by a major technological company rather than SenseMe. This difference could be attributed to the fact that the L group had participated in a live deployment and had first hand experience with SenseMe.

Saving	Group L		Group O	
	SenseMe	Similar app	SenseMe	Similar app
1% salary	μ = 3, σ = 1.89	μ = 3.9,σ = 1.66	μ = 4.36,σ = 1.95	μ = 3.95, σ = 1.98
10% salary	μ = 2.4, σ = 1.58	μ = 3.1, σ = 1.97	μ = 3.47, σ = 1.96	μ = 3.4, σ = 2.01
10 minutes	μ = 3.2, σ = 1.81	μ = 4, σ = 1.76	μ = 4.62, σ = 2.1	μ = 4.14, σ = 2.18
1 hour	μ = 2.1, σ = 1.1	μ = 2.9, σ = 1.73	μ = 3.58, σ = 1.95	μ = 3.49 , σ = 2.0

Table 3. Relationship between participants' group and responses to the Brand Reputation questions

4. Design Guidelines

The aim of our study has been to explore privacy, trust, risks involved and other related issues in smart phone based personal sensing systems and applications. In this section, we propose some high-level guidelines that we have derived from the results presented and which we believe would be useful for personal sensing app or system designers in order to mitigate these concerns. We also outline some interesting directions for future research.

4.1. Maintaining Transparency

As evident from our results, many users are wary of apps that use their data for purposes other than what it declares. Also, they do not want apps to have unauthorized access to any data, sensors or services. Hence, its important that the app or system designer maintain transparency in the design and documentation of the app. They should clearly define what the purpose of the app is, and what sensors, content and services it will access. Moreover, this information should be provided to the users at the time of app install or download and via a medium that would be most pervasive and accessible. As evident from the results, the EULA may not be the best medium since 50% of the users never read it. Thus, a medium other than the EULA (say, a dedicated pop-up message that catches the users' attention at the time the app is installed) may be more suitable for this purpose.

In addition, further research should be conducted to test the best formats and media to educate users about an app's use of private data. Previous research [11] has demonstrated that displaying required permissions and privacy information in a clearer fashion could play a more active role in influencing users to make privacy protecting decisions at the time of app selection. Lin et al. [12] found that informing users about the purpose of an app's access to phone resources improved decisions and eased privacy concerns. Research has been carried out on online privacy notices formats for websites [13, 14]. Similar studies on privacy notices for mobile apps will be beneficial taking into consideration the mobile devices form factor, user attention span, user demographics etc.

4.2. Access to sensitive data or intrusive sensors

As evident from our results, if an app or system requires access to highly sensitive data such as photos, contacts or health data or to intrusive sensors such as GPS or microphone, many of the users are reluctant to use it unless they have very specific needs for it. There is a significant fraction of users (33%) who feel very strongly about this and would never use an app that accessed such data or sensors. Hence, it is imperative that the designers avoid this and find other alternatives. For instance, rather than using raw GPS coordinates, it could be sufficient for an app to know which general geographic zone the user is currently in if accuracy is not a major concern. Or the designers could design the app to sense a user's location via network or cellular provider to localize the user at a city or locality level, instead of using GPS to localize the user to an exact location.

However, a significant number of popular smartphone apps such as Foursquare and Facebook make use of location. These popular commercial location sharing apps seem to mitigate users' privacy concerns by allowing them to selectively report their location using check-in functionalities instead of tracking them continuously and automatically. Hence, selective sharing of accurate location, sensed from GPS, based on the users' discretion is also a viable option.

4.3. Sharing of the sensed information

Sharing users' data with other applications, other users of the same app or with their friends on a social network should be done at the users' discretion. Thus, the designer should give the users' full control of what they want to share, when they want to share it and if they want to opt out of this facility. Alternatively, developers could provide annotations that reflect their privacy and data sharing policies, and this information could be incorporated into warnings or data access requests. Similar to tools like AppFence [15], that tell users whether their data is being sent to advertisers or other known third parties, other tools should be developed that can inform users if their data is being shared with other applications, other users or being posted publicly anywhere.

4.4. On-device vs on-cloud/server

Since most users do not want their data to be stored in the cloud or on a third party server, designers of the system or app should consider the design in such a way that majority of the processing and storage is carried out on the user's device. The subjects' responses suggest that they would prefer that only a limited amount of data is transmitted over the network in a secure or encrypted manner and stored on a cloud or server. In addition, the system or app designers should clearly declare to users, what data is being processed on the device and what data will be transmitted to a server, and how it will increase the benefit or value of the system to the users. If an app is accessing and transmitting more data than it has declared in order to provide the required services, it can be easily discovered by users who use the app, and this could potentially be reflected in the app's reviews.

While Balebako et al. [16] have experimented with notifying the user and visualizing the amount and type of information being shared on his device in order to understand his perceived concerns, they do not comment on whether indicating the benefit or value of the shared information would influence the users' decision to share it.

4.5. Data for Benefits

As evident from our study, users are more willing to share their data if the benefits to them are significant, such as saving them time or money, or if the app

provides them with services such as timely and relevant information. Though our study does not determine what payoff would result in an optimal number of users sharing private information, it is evident from the results that most users place a value on their private information which dictates how and under what circumstances they would surrender it. Thus, the app or system designer should balance privacy invasion with the benefits to users in a way to maximize user participation. Staiano et al. [6] have investigated the monetary value that people assign to different kinds of personal information as collected by their mobile phone, including location and communication information.

4.6. Usage of the sensed information

Our study demonstrates that if the users' sensed information is utilized for a beneficial and effective cause such as saving lives, as opposed to commercial purposes and targeted advertising, users' are more willing to share it. Yet this increased willingness to share sensitive data is not without reservation, as evident by the fact that more than 25% of the subjects were unwilling to share such data even if the benefits of sharing data can result in saving lives. This may be attributed to the negative connotation that surrounds usage of such data for commercial and advertising means. One subject opined that the only use of sensitive data by a sensing party would be to capitalize on it for targeted advertising. Another stated that he was skeptical of the inability of companies to not monetize on the collected data, making him very unwilling to share sensitive data even if the intention of the collecting party was for a good cause such as saving lives. Nonetheless, it is essential that the app designer clearly state their intentions for the usage of the sensed data.

4.7. Build trust and reputation

As we observed, users are more willing to trust a known brand or company with their data as opposed to an unknown entity. Hence, it is important to establish a reputation and gain their trust. As evident from our results, subjects were more likely to use an app created by a well known brand then a smaller, obscure company. Yet if they have first hand experience with the app, they are more willing to use it. Additionally, data that they claim, they would never share with a cell phone application, do not align with what they currently share with major email providers. All but two respondents use an email service provided by a well known brand, that uses email content and contact listings, in order to enrich the email experience; something that one-fourth of the users claimed they would never allow on a cell phone app.

Moreover, we believe that users' reluctance to share their data, even if to save others' lives, also relates to trust. This reluctance of an individual to provide resources that can save another's life and require no effort on his/her part, should be startling to app designers. This should be a testament to both the lack of trust and the importance to regaining this trust.

Hence, one of the ways via which designers can build trust is to perform live deployment of their systems and apps in the wild. This could be conducted with the aid of professional market research companies which recruit users from different demographics. The designers can conduct these deployments in-lab and *in situ* with follow-on user feedback surveys. This would allow the users to gain first hand experience with their apps in a practical real life scenario and allow the designers to get valuable feedback on their apps.

5. Limitations

In what follows we discuss some limitations of this work.

The Likert scale ratings and open-ended questions utilized in our studies are not absolute measures of user concern because our surveys explicitly asked respondents about privacy, trust and their willingness to use these apps. Surveys that directly ask questions about privacy may suffer from inflated user concerns about privacy [10] and therefore are not reliable measures of absolute levels of concern. We expect that this applies to our study as well.

Our survey questions compared users' responses to different mechanisms and alternatives for data sharing, data storage and retention, data collection, and benefits in terms of time and money etc. Where appropriate, we provided users with scenarios, which were meant to help them with assessment of possible risks. In our results, we weighed the risks and involved trade offs relative to each other. Thus, the same set of priming biases are applied equally to all of the alternatives presented in the surveys, so the priming effect should not influence the results.

We do not claim to predict the users' decisions when confronted with these risks and trade-offs in real life because our study relies on self-reported data. As with the priming bias, we do not believe that self-reporting affects the validity of our results because this bias is equally present for all alternatives.

Our web-based survey did not reach professionals from various other backgrounds, who may have different concerns. However, it was taken by a large number of participants with varying ages, demographies and occupations. Secondary studies may be needed to target specific groups that could potentially have their own privacy and security concerns, such as doctors (who handle health records), lawyers (who handle client

data), or company executives (who handle corporate data).

6. Related Work

Since we explore a multitude of privacy, trust, risks involved and other related concerns with personal sensing systems and applications, we have divided the related work into various sections and differentiate our work from them. Moreover, as stated before, none of the existing works have addressed such a broad spectrum of concerns with personal sensing applications.

6.1. Privacy concerns with personal sensing via proprietary devices

Klasnja et. al. [17] explored privacy concerns with personal sensing in a field trial of the UbiFit system[18]. Unlike SenseMe, which runs on the user's smartphone, the sensing in UbiFit was carried out using proprietary hardware and hence, wasn't privy to personal information. For this reason, our user study focused more on what was being sensed and inferred as opposed to the sensors being used. Also, UbiFit only recognized physical activities while SenseMe performs temporal context and activity recognition along several dimensions. Moreover, we investigate several factors such as brand trust, recognition and awareness in addition to just sensors. However, our results support their arguments on Data Retention and Perceived Value of the applications.

6.2. Privacy concerns with location tracking and sharing

Mobile privacy research [19, 20] has traditionally focused on location tracking and sharing and has examined users' privacy concerns about sharing mobile location data. Iachello and Abowd [21] described how to build appropriate privacy controls into a social location-sharing application.

Numerous studies [22–24] have explored location sharing behavior of users and the factors that influence it. Their findings suggest that who is requesting the user's location, why they are requesting it and to what level of granularity significantly affects the user's decision to share it. Lederer et al. [25] found that the identity of the location requester matters more than the place in a user's willingness to share his or her location, while Anthony et al [26] focused on the effect of the specific place that the user is asked to share. Moreover, it seems that the users' age, gender, mobility, and geographic region also play a role in location sharing behavior. While we asked users in general about sharing various forms of data in addition to location (such as activities), similar factors may influence users' decisions. In addition, we questioned participants about their privacy concerns related to all types of fine-grained personal data from their smartphone. Hence, it is difficult to directly compare our work with them.

6.3. Other smartphone privacy concerns

Smartphone apps have the ability to access a number of resources beyond location data. Smartphone APIs let applications read many types of data (e.g., photographs) and make changes to the phone (e.g., delete data). Few studies have explored the space of smartphone privacy and security beyond location.

Lane et. al. [27] discuss the issue of privacy in their survey of smart phone sensing applications and systems. However, they do not present any evaluations or quantifiable results. Instead they draw conclusions from existing work in smart phone sensing.

Muslukhov et al. [28] asked 22 smartphone users about the value and sensitivity they assigned to eleven types of data (SMS messages, photos, contacts, emails etc.) on their phones. One aspect of our study, where we asked users about Data Control (see Section 3.3) and list the types of information for which they would limit or disallow sensing, shares similar goals with this work.

Felt et al. [29] report that users' concerns about data sharing depend on who the data is being shared with. Their findings suggest that for different data types, publicly sharing the data most concerning than sending the data to a server. Sharing with friends and advertisers rank in the middle, between public sharing and sending the data to a server. While we asked users about data sharing as well (Section 3.4), the sharing mechanisms that we suggested were different.

Staiano et. al. [6] performed auctions of users' data and found an optimal price for which users would sell it. This auction however released the data in general and not to a commercial entity looking for ways to capitalize on the data and target the users. Our results show that users are not uniform in their sharing behavior or how they feel their data should be used. Thus varying the consumer of the data might greatly influence a user's ability to share it.

7. Conclusion

In this paper, we conducted an exploratory study of privacy, trust, risks and other related concerns of users with smart phone based context-aware personal sensing applications. We reported results obtained from a live deployment with a smart phone sensing application and a web-based study involving 70 participants in all. Our results show that users are concerned about their sensed data being misused or used for commercial purposes. They are also concerned that the app or system may have unauthorized access to sensitive content. Also, the users want more control of what data they want to share, when they want to share it and

with whom. However, they are willing to trade privacy for additional significant benefits such as saved money and time or if their sensed information is used for beneficial and effective causes such as saving lives. In addition, they are willing to trust reputed technology companies which have a brand name, if the benefits are significant, despite being aware that their data is sensed and collected by these companies.

Based on these results, we proposed a few design guidelines for smart phone based personal sensing system or app designers. These include maintaining transparency in the design and documentation of the app, using alternative solutions instead of accessing sensitive data or intrusive sensors, as well as sharing the users' data at their discretion. Moreover, designers should design the app to carry out majority of the storage and processing on the device with only a limited amount of data being transmitted over the network. The designers should also balance privacy invasion with additional benefits to users and establish trust and reputation by performing live deployments of their systems in the wild.

In future, research maybe required to determine ways to remove the stigma surrounding the possibility of selling of sensitive information by companies. In addition, persuasive computing techniques should be researched to persuade users who are hesitant to use sensing apps and share their data even when the payoff is something as powerful as saving lives.

References

[1] P. Bhargava, N. Gramsky, and A. Agrawala, "Senseme: a system for continuous, on-device, and multi-dimensional context and activity recognition," in *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2014.

[2] J. Krumm and R. Hariharan, "Tempio: inside/outside classification with temperature," in *Second International Workshop on Man-Machine Symbiotic Systems*, 2004.

[3] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell, "The jigsaw continuous sensing engine for mobile phone applications," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, 2010.

[4] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell, "Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*, 2008.

[5] P. Zhou, Y. Zheng, Z. Li, M. Li, and G. Shen, "Iodetector: A generic service for indoor outdoor detection," *Proceedings of the 10th ACM Conference on Embedded Networked Sensor Systems*, 2012.

[6] J. Staiano, N. Oliver, B. Lepri, R. de Oliveira, M. Caraviello, and N. Sebe, "Money walks: a human-centric study on the economics of personal mobile data," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2014.

[7] N. S. Good, J. Grossklags, D. K. Mulligan, and J. A. Konstan, "Noticing notice: a large-scale experiment on the timing of software license agreements," in *Proceedings of the ACM SIGCHI conference on Human factors in computing systems*, 2007.

[8] Y. Rogers, H. Sharp, and J. Preece, *Interaction design: beyond human-computer interaction*. John Wiley & Sons, 2011.

[9] M. Borenstein, L. V. Hedges, J. P. Higgins, and H. R. Rothstein, *Introduction to meta-analysis*. John Wiley & Sons, 2011.

[10] A. Braunstein, L. Granka, and J. Staddon, "Indirect content privacy surveys: measuring privacy without asking about it," in *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 2011, p. 15.

[11] P. G. Kelley, L. F. Cranor, and N. Sadeh, "Privacy as part of the app decision-making process," in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 2013.

[12] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 2012, pp. 501–510.

[13] A. M. Mcdonald, R. W. Reeder, P. G. Kelley, and L. F. Cranor, "A comparative study of online privacy policies and formats," in *Privacy enhancing technologies*. Springer, 2009, pp. 37–55.

[14] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor, "Standardizing privacy notices: an online study of the nutrition label approach," in *Proceedings of the ACM SIGCHI Conference on Human factors in Computing Systems*, 2010.

[15] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall, "These aren't the droids you're looking for: retrofitting android to protect data from imperious applications," in *Proceedings of the 18th ACM conference on Computer and communications security*. ACM, 2011, pp. 639–652.

[16] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen, "Little brothers watching you: Raising awareness of data leaks on smartphones," in *Proceedings of the Ninth ACM Symposium on Usable Privacy and Security*, 2013.

[17] P. Klasnja, S. Consolvo, T. Choudhury, R. Beckwith, and J. Hightower, "Exploring privacy concerns about personal sensing," in *Pervasive Computing*. Springer, 2009, pp. 176–183.

[18] S. Consolvo, P. Klasnja, D. W. McDonald, D. Avrahami, J. Froehlich, L. LeGrand, R. Libby, K. Mosher, and J. A. Landay, "Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays," in *Proceedings of the 10th ACM international conference on Ubiquitous computing*, 2008.

[19] D. Cvrcek, M. Kumpost, V. Matyas, and G. Danezis, "A study on the value of location privacy," in *Proceedings of the 5th ACM workshop on Privacy in electronic society*, 2006.

[20] G. Danezis, S. Lewis, and R. J. Anderson, "How much is location privacy worth?" in *WEIS*, vol. 5. Citeseer, 2005.

[21] G. Iachello and G. D. Abowd, "From privacy methods to a privacy toolbox: Evaluation shows that heuristics are complementary," *ACM Transactions on Computer-Human Interaction*, vol. 15, no. 2, p. 8, 2008.

[22] S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge, "Location disclosure to social relations: why, when, & what people want to share," in *Proceedings of the ACM SIGCHI conference on Human factors in computing systems*, 2005.

[23] N. Li and G. Chen, "Sharing location in online social networks," *Network, IEEE*, vol. 24, no. 5, pp. 20–25, 2010.

[24] D. Wagner, M. Lopez, A. Doria, I. Pavlyshak, V. Kostakos, I. Oakley, and T. Spiliotopoulos, "Hide and seek: location sharing practices with social media," in *Proceedings of the 12th ACM international conference on Human computer interaction with mobile devices and services*, 2010.

[25] S. Lederer, J. Mankoff, and A. K. Dey, "Who wants to know what when? privacy preference determinants in ubiquitous computing," in *CHI'03 extended abstracts on Human factors in computing systems*. ACM, 2003, pp. 724–725.

[26] D. Anthony, T. Henderson, and D. Kotz, "Privacy in location-aware computing environments," *IEEE Pervasive Computing*, vol. 6, no. 4, pp. 64–72, 2007.

[27] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell, "A survey of mobile phone sensing," *Communications Magazine, IEEE*, vol. 48, no. 9, pp. 140–150, 2010.

[28] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov, "Understanding users' requirements for data protection in smartphones," in *Data Engineering Workshops (ICDEW), 2012 IEEE 28th International Conference on*. IEEE, 2012, pp. 228–235.

[29] A. P. Felt, S. Egelman, and D. Wagner, "I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns," in *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*. ACM, 2012, pp. 33–44.