

Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use

Irina Shklovski¹, Scott D. Mainwaring², Halla Hrunn Skúladóttir¹ and Höskuldur Borgthorsson¹

¹ITU Copenhagen

Rued Langgaards Vej 7

2300 Copenhagen, S. Denmark

irsh@itu.dk, {hallahsk,hoskborg}@gmail.com

²Unaffiliated

Portland, Oregon, USA

scottmainwaring@gmail.com

ABSTRACT

Mobile devices are playing an increasingly intimate role in everyday life. However, users can be surprised when informed of the data collection and distribution activities of apps they install. We report on two studies of smartphone users in western European countries, in which users were confronted with app behaviors and their reactions assessed. Users felt their personal space had been violated in “creepy” ways. Using Altman’s notions of personal space and territoriality, and Nissenbaum’s theory of contextual integrity, we account for these emotional reactions and suggest that they point to important underlying issues, even when users continue using apps they find creepy.

Author Keywords

Mobile devices, data privacy, bodily integrity, learned helplessness

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

“I feel like I am constantly tracked either on my phone or on my computer, at stores, etc. I feel like this is a part of daily life now” – Survey, USA

What does it mean to “feel” like you are constantly tracked by unseen and largely unknown entities? In the quote above who was doing the tracking was not named and for most of our respondents through the interviews and the survey they remained unknown, largely abstract entities that somehow collect personal, sometimes even intimate information. With the rise of personal computing in the 1980s and especially of the world-wide web in the 1990s and 2000s, and the movement of computers out of offices and into homes,

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2014, April 26 - May 01 2014, Toronto, ON, Canada
Copyright 2014 ACM 978-1-4503-2473-1/14/04...\$15.00.

<http://dx.doi.org/10.1145/2556288.2557421>.

computing became understood as technology of mass consumption and marketing, raising new questions about consumer tracking (especially via web cookies) and targeting [24]. Mass data collection is becoming “designed in” to everyday life [36] as we increasingly rely on technologies capable of monitoring, storing and distributing information about us. Occasionally data collection by governments, organizations, corporations or even small application developers comes to light causing a public outcry [39]. Yet much of the tracking and data collection does not necessarily happen with insidious goals, but is a result of businesses attempting to deliver personalized services and advertising more effectively [40]. Despite potentially positive outcomes of this kind of tracking, the public response remains negative often because people simply find it *creepy*. At the same time, people continue to use the technologies and applications implicated in data collection even as they express outrage over such activities. This is typically seen as an expression of the “privacy paradox” where intentions and behaviors around information disclosure often radically differ [2,32,40].

Over the last decade, especially since the launch of the first app store by Apple in 2008, computing has moved out of the home and into our pockets, purses, hands, and even bedsides, bringing with it a new wave of concerns, increasingly framed in terms of “creepiness.” We will define *creepy* more carefully later on in the paper but we note here that we use this term to denote an emotional response to a sense of wrongness that is difficult to clearly articulate. We argue that the notion of creepiness is a concept worth unpacking, and by way of example analyze a corpus of data recently collected from smartphone app users in a number of western countries using perspectives of bodily integrity, personal space and territoriality [3,30,31,34].

From this analysis a different notion of “users” emerges. Listening closely and empathically to our study participants, more than seeing them just as decision-making agents, we began to identify them as people¹ caught up in a

¹ For the purposes of this paper, when we say “people” we are referring to fairly affluent consumers in developed Western countries. Questions of personal space and privacy are clearly culturally dependent, and from, e.g., Chinese or Global South perspectives smartphones may have very different meanings.

bind over the limits of their personal space: On the one hand, smartphones felt like intimate zones and extensions of their bodies. On the other hand, in order to empower themselves through this boundary-enlarging technology, they felt they had to accept the possibility of unknown, foreign access into this zone of expanded intimacy. Reframing the “privacy paradox” in these terms opens different opportunities and responsibilities for technology development and policy design.

SMARTPHONES AS EXTENSIONS OF THE SELF

Ubiquitous, nearly infinitely personalizable and most often individually owned at least in the Western world, mobile phones have variously been discussed as deeply personal extensions of the self [26,29] or even organic parts of the self [23,33]. As mobile technologies developed and become domesticated, smartphones have moved from an expensive curiosity to a commodity deeply integrated into their users' everyday lives. The functions these devices perform are not just about practical problem solving, escapist entertainment or convenient extensions of sociability, but also about projecting and constructing the self [12].

Like our homes, smartphones are implicated in the way of being and deeply embedded in the life experience. Yet in HCI research aesthetic or emotional experiences are less of a concern than the serious business of sociability, friendship, collaboration and getting ahead in life. Finding a particular experience “creepy” is a kind of emotional experience, but it can also serve as an indication of a disjuncture between the way developers and users experience and interpret functions or applications. Much of the discussion of discomfort or feelings of creepiness in the course of technology use in HCI has been conflated with user privacy concerns of data access and security. We use two theoretical frameworks familiar to HCI researchers for thinking about privacy, but extend the conversation to consider how creepiness might be conceptualized in this context.

Altman's notions of personal space and territoriality

Irwin Altman's theory of privacy regulation has been discussed in prior work [34]. Altman conceptualizes privacy as a process through which people attempt to achieve a desired privacy level in any situation in life by selectively adjusting access to themselves through controlling information they disclose or receive [3]. In other words, people keep secrets and manage themselves and their secrets through ongoing acts of revealing and concealment [30]. Altman conceptualized privacy as one of the four key concepts central to the study of environment and social behavior: privacy, personal space, territoriality and crowding. Altman explains personal space and territorial behavior as mechanisms people use in the service of privacy goals to regulate interpersonal boundaries and environmental factors.

Personal space is typically conceptualized as a “boundary around a person, intrusion into which is often uncomfortable and generally not permitted” [3]. Intrusions or invasions

of personal space represent privacy violations and can result in a variety of responses from feelings of anxiety to actions designed to increase distance and reduce interaction.

Territorial behavior is another type of a social regulation mechanism. Altman argues that just like animals are territorial about their feeding or mating areas, people also regard certain places and objects as their personal or primary territories. A primary territory usually refers to an area, a place or an object that fulfills certain needs or motives and where ownership is clearly conveyed through some form of personalization. Entry into another person's primary territory is typically done with permission only and trespassing is considered an invasion that can elicit intense emotional responses and physical action [3].

Several researchers have explored user decisions to allow physical access to their smartphones and personal data, showing that people regard their phone as highly private and are reluctant to share it with others, expressing discomfort even when sharing it with close friends [16,17]. Research has also shown that smartphone users are surprised and feel violated when they find out that applications are accessing data on their smartphones ostensibly without their knowledge [10,25]. In other words they feel that this data access breaches their privacy norms and should be done with their informed consent, the same reaction people show when their primary territory is breached.

Increased computerization affects individuals and society as a whole as it brings increased imbalance in power between individuals and large institutions. In the case of smartphones, the power balance between the smartphone user and the application developers is quite uneven. The user has few means of safeguarding his territory and often she is unaware of the encroachment. Her options include shifting the desired level of privacy towards the actual level; a common approach used when repeated attempts to regulate privacy have failed.

Nissenbaum's contextual integrity model

Helen Nissenbaum's theory of contextual integrity complements Altman's notions of personal space by calling attention to context-dependent social norms and regulated information flows [31]. Even from within the most private, personal space, people never interact with information in the abstract, but always as it is embedded in a particular social context. Nissenbaum defines contexts as “structured social settings characterized by canonical activities, roles, relationships, power structures, norms (or rules), and internal values (goals, ends, purposes)” [31]. Smartphones may be extensions of personal space, in Altman's terms, but through them people participate in a variety of quite different contexts. These may indeed vary on an app-by-app basis: what is socially legitimate information sharing for a health app might include highly sensitive information about one's body and behaviors, whereas what it is socially legitimate for a game like Fruit Ninja to share, while difficult to

precisely define and probably not yet worked out by its many stakeholders, is sure to be much less sensitive.

Contextual integrity is maintained if information flows according to contextual norms. However, digital technologies make it easy to disregard such norms either through ignorance or intentionally [4]. That a digital system is respecting contextual integrity may be difficult and expensive to attain, often beyond the capabilities of any particular individual. However, individuals can sometimes (and dramatically) detect when a breach of norms has occurred when they see that information they consider off limits with respect to the context in which it was given is being used in ways antithetical to its socially sanctioned expectations.

What are we feeling when we are creeped out?

Encountering a violation of contextual integrity of one's information can result in a range of emotional responses, though it may not necessarily lead to outward action to alleviate the problem. One can be outraged, exasperated, horrified, even cynically bemused. Along with or in addition to these, people often describe their discomfort by referencing the word "creepy".

Collins English dictionary defines *creepy* as "having or causing a sensation of repulsion, horror, or fear, as of creatures crawling on the skin" [8]. The reference to creepy-crawlies on the skin grounds this experience in a certain kind of fundamental, embodied boundary violation, in Altman's terms. In Nissenbaum's terms, creepy information flows often involve realizations that personal secrets have been, or could be, revealed to those who have not been explicitly granted access to them. "Creepy" does not necessarily signal that the end result (contextually appropriate advertisement for example) is not beneficial to the user, but to a kind of dual violation of deeply held expectations not just of "contextual integrity" in terms of respecting social norms, but "personal zone integrity" in terms of control over sensitive boundaries.

TRACKING AND DATA LEAKAGE

Discussions of information disclosure practices via mobile phones have been extensive though mostly focused on location-based services and interpersonal information disclosure [22,28]. Researchers have turned their attention to smartphone app data distribution, often termed "data leakage" to third parties without direct owner notification only recently [25]. Ostensibly, users give permission for smartphone applications to access their use data and to utilize it for a range of purposes when they press "agree" on the End User License Agreement (EULA) screen in the process of installing a smartphone application. Despite this, revelations of data distribution by many common applications to third party advertisers have resulted in expressions of concern and discomfort by smartphone users [9,10,19,25].

Where does the data go?

In a survey of 5000 mobile users reported by Think Insights, 82% of the participants claimed to notice ads on their smartphones and 49% had taken action based on those ads². Advertising in applications has a substantial effect on users and their purchase patterns. The global mobile advertising market was valued around 5.3 billion dollars in the year 2012 illustrating just how valuable it is for advertisers to have access to user patterns and personal data [14]. Ad libraries and analytic companies buy the data that smartphone apps collect, analyzing and aggregating user data to build a profile for targeted advertising.

Imagine a smartphone owner who has on her phone a typical flashlight application like the High-Powered Flashlight from iHandy inc., and a simple one-player game application like Fruit Ninja from Halfbrick Studios. Both apps require access to the phone number and device ID, the remote number if a call is active and which other applications are used on the device. Additionally the flashlight app has access to precise location, hardware controls and system tools that enable it to change system configuration and display settings. The Fruit Ninja app also requires access to the list of accounts known by the device. The information these apps collect is sold to several third parties, one of which is Flurry, a big ad and analytics company that according to its website prides itself on taking in two terabytes of data from 2.8 billion app sessions per day [11]. Flurry buys data from both these apps and can build a profile of a user that includes information about what type of phone she has, what applications she uses and how frequently, what accounts she has on her phone, which numbers she calls (which can link to other consumer profiles in the Flurry database) and precise location of the phone which, if logged with timestamps, gives an accurate estimate of where she lives, works and travels. By properly aggregating this information the company gets a detailed profile of her as a consumer, information that is worth money.

Smartphone users' response to tracking and data collection as well as the willingness to share data with others varies but users are invariably surprised at the extent of data leakage on their phones [5,22,25]. Although many are willing to share data with developers in the right context, particularly when the type of information collected and its purpose is made clear, it is difficult to justify why gaming applications that are about fruit slicing or sling shooting birds would need the user's location, phone number or information about other running apps. The user in the example above has given her permission for the Flashlight and Fruit Ninja apps to collect and distribute her information at the point of installation when she clicked through the End User License Agreement (EULA). Yet smartphone users often have little

² Think Insights by Google reporting on a study conducted by Ipsos OTX. Retrieved from <http://goo.gl/SRTh4>

³ In our study we used <http://privacyscore.com/mobilescan> and

understanding of which permissions they consent to and what these mean. Moreover, studies show that most users simply ignore permissions completely [9,10]. People ignore EULAs for various reasons, for example warning fatigue [10], lack of motivation [15], or an overriding desire to install the app where reading the EULA is seen as a cost without a clear benefit [25].

“Data leakage” at one level appears to be a misnomer — the data flows the app is enabling are most likely not accidental oversights or unintended consequences, but central to the business model that makes the app possible in the first place. “Intended data distribution” would seem more apt a term and this is how tracking is often included in EULA. However, upon closer inspection, calling this *leakage* captures an important aspect of the consequences of installing the app: if the user conceptualized her personal space as having integrity, and a more-or-less controllable border, from her perspective the app makes this border “leaky”, and intentionally so.

While ensuring readability and usability of EULA is important, here we consider why people continue to use apps even when they know these apps reveal personal data to advertisers. We also discuss why it is important to consider these issues both at the individual level and in the broader social context. The studies we present consider what it means to lose control over access to the self to unknown third parties with few concrete or obvious consequences. Why does it matter that third-party advertisers can gain access to certain types of personal information via smartphones? How do people react to finding this out and what might they be willing to do about it?

STUDYING ATTITUDES TOWARD DATA LEAKAGE

We conducted two studies investigating how smartphone users feel about data access on their phones. The first study was a qualitative investigation conducted in two countries in order to assess whether smartphone users may be willing to change their behavior once they have been informed about tracking and data leakage. We exposed our participants to data leakage to see if they made changes in their user patterns when informed of potential privacy breaches. In the second study we conducted a survey on an international opportunity sample where we explored how data privacy sentiments might affect smartphone users when choosing and downloading applications. Furthermore we explored how important it was for our participants to maintain their privacy while using applications on smartphones.

Study 1 — Interviews

In February 2013 we conducted 13 semi-structured interviews in two countries: 6 in Iceland 7 in Denmark. We recruited users from our own personal networks and by advertising on Facebook. The prerequisite was that our users owned and used Android smartphones and were between 25 and 60 years of age. Our sample consisted of 8 men and 5

women, and the age distribution was between 27 and 55 years with a mean of 33.5.

We chose to use the Android OS as our research platform because the system presents the user with the permission screen each time she downloads an application. Our participants then would have encountered the terms of use more often than an iOS smartphone user.

The interviews were conducted in English, Danish or Icelandic depending on participant preference and lasted 40-70 minutes each. We started with general questions regarding smartphones and applications and then asked participants to think aloud while demonstrating how they usually go about accessing the Google Play Store on their mobile phones and downloading an application of their choosing. After our participants downloaded an app, we moved on to questions regarding attitudes and beliefs about data privacy. We also asked users what they think their apps do with the data they access, to get insight into participants’ mental models of data security on smartphones.

In the second half of the interview we introduced some of the most common permission statements found on popular apps in Google Play and discussed what they mean. We also showed them a website that allows users to scan apps, and asked users to download an app in Google Play that scans permissions and analytic libraries connected to apps currently located on the user’s mobile phone³. Finally, we asked participants to download and try out a popular game on Google Play, Fruit Ninja, then explained its main functions and discussed the permissions our users had accepted before downloading the app to their mobile devices.

Three weeks after the initial interviews we sent our participants a follow-up questionnaire where we asked them questions regarding their use patterns and conducted a follow up interview via email or phone. All interviews were transcribed and coded. We used an iterative analysis process of open and thematic coding throughout the data collection process. This allowed us to test emergent themes in follow up interviews.

Study 2 — Survey

In March and April of 2013 we deployed a survey to gain a better understanding of privacy sentiments of smartphone users more broadly. We obtained an opportunity sample by recruiting via our own networks on Facebook and by sending participation requests to university student mailing lists in Iceland and Denmark. A total of 272 respondents answered our survey and 187 completed it (116 women and 71 men). Although this was an opportunity sample it was quite diverse. The majority of participants were between the

³ In our study we used <http://privacyscore.com/mobilescan> and Addons Detector by denper available from the Google Play store <http://goo.gl/TM9g2>

ages of 25 and 44 (81%) although participants ranged in age from 18 to 74. Over 60% were highly educated (graduate degree) and approximately 13% had less than a bachelor degree. Most of the respondents lived in Denmark (43%), Iceland (20%), USA (20%) and Sweden (10%). However, the sample included participants from 12 countries. Although slightly over half of our respondents were veteran smartphone users (51%), having used a smartphone for over 2 years, 18% were new to the smartphone world, having used the device for less than a year.

We measured attitudes toward app stores, personalized advertising and general sentiment toward data collection by businesses as well as the general level of concern about threats to personal privacy. We also developed a set of questions focused specifically on attitudes towards data collection and management by smartphone apps.

In order to measure attitudes toward data collection by mobile applications respondents were randomly presented with either Fruit Ninja game app or Flashlight app with the developer description of the app. Below the description we listed permissions that the app requires the user to agree to before downloading and asked respondents to indicate which they felt permissions were appropriate or inappropriate given the functionality of the app.

On the next page we explained which ad-libraries and analytics companies received the data the apps collected from the mobile phones and added the explanation these companies generally give for how this data is utilized. We followed up with questions regarding awareness about data collection and distribution in mobile applications and whether participants read the EULA prior to downloading apps. The final question of the survey was free response, asking users to share how they felt about data distribution to third parties in general. Due to space constraints all survey instruments are available upon request.

Nearly 50% (89 of 187) of participants that had completed the survey responded to this question and we coded these responses using an iterative coding process by moving from an initial open coding through axial and thematic coding to distill thematic categories. Where necessary, quotes presented in this paper were translated from Icelandic or Danish. Quotes and excerpts are marked Interview or Survey to indicate the source and Country to indicate location of the respondent.

FINDINGS

In broad strokes, many of our findings confirmed other studies of attitudes toward mobile phone data leakage and privacy. Similar to the findings in [7] our interview participants became more concerned with data leakage as a kind of privacy violation throughout the interview. Similar to findings in [25] they also indicated that these privacy concerns were often overridden by other factors when they made decisions about whether to install a particular application. As in [9,10,17] our participants had strong opinions on

what data they considered sensitive (emails, images and messages) and substantial misconceptions about what data applications they used were accessing on their mobile phone, as has been noted by [10,25].

Just like the participants in [10] and [15] our interviewees ignored the EULA when they were asked to install an application of their own choosing from the app store during the interview. Yet they had typically expressed concern about their privacy in the prior portion of the interview, discussing privacy implications of smartphone application data leakage. When asked why they had ignored the EULA two explanations dominated similar to those described in [9,18,25]. First, they had never had any real negative consequences from data collection. Second, the desire to have the application trumped any concerns for data collection, which was explained away as “the way things are.”

“Yes, I used to think about it, but I decided to stop it. I know I am agreeing to this (data collection). If I really want the app, I take my chances.” – interview, Denmark

Other researchers have argued that it is important to redesign the EULA to help users be more aware of what is going on in their smartphones [25,41]. In the final portion of the interview, participants were asked to install Fruit Ninja and researchers walked them through the EULA explaining all permissions requested by the application. Similar to [25], this exercise revealed a mismatch between the interviewees’ mental models of what Fruit Ninja might reasonably request in order to function and the kind of data it actually requested. That is, they expected the data collected via the app to be contextually appropriate to a single-player game of slicing fruit. Not surprisingly, this mismatch caused participants significant discomfort similar to [9,24]:

“I feel deceived, I had no idea it was this systematic” – interview, Iceland

During the follow-up interview three weeks later our participants said they wanted to be more cautious and were trying to pay attention to what they were downloading and installing on their mobile phones. More than half (7 of 13) had deleted an application although in most cases it was Fruit Ninja:

“I deleted Fruit Ninja because it was gathering a lot of information, although it might also be the case with other apps on the phone I don't want to delete them, I simply enjoy them too much.” – interview, Iceland

This was to be expected as it was not their choice to download Fruit Ninja and the majority of our participants did not have prior experience with this app. Even though participants were not willing to trade personal data for useless applications they accepted the possibility that other more “meaningful” applications might be doing the same thing. Although some approached their smartphones with more caution they admitted they had not changed their practices very much:

“I just get the applications I want, I had forgotten all about this data gathering stuff.” – interview, Denmark

Similar to findings in [15] and in slight contradiction to [18] the extensive explanations and demonstrations of what can be done to monitor privacy more closely, installation of detector applications on participants’ phones and introductions of trackers websites, our participants did not use these tools, or looked for new ones, because they regarded it as a waste of time:

“I will never waste my time reading privacy policies, my time is simply too valuable. To some extent I just have to accept this.” – interview, Iceland

The survey findings followed a similar pattern. Our participants were generally aware of data collection and distribution to third parties and believed that more data was collected than necessary. Nevertheless, few respondents restricted mobile data while using their phones. Indeed, 78% had data turned on continuously with 56% saying they use data a lot. Similar to findings in [10] nearly 57% said they had deleted an app and roughly 62% had aborted the installation process of an application because of privacy concerns.

Overall survey respondents were concerned about information disclosure to businesses (mean=2.29, sd=.487, range 0-3 where 3=high concern) and reported moderate levels of concern regarding threats to personal privacy (mean=2.01, sd=.466, range 0-3 where 3=high concern). Participants were also quite concerned about smartphone integrity and their privacy in regards to that, which manifested in low scores on our scale of how trusting smartphone users were toward the integrity of app stores and app developers (mean=1.23, sd=.65, range 0-3 where 3=highly trusting).

Beyond the Privacy Paradox

Results from both studies clearly illustrated what is commonly called a privacy paradox where our participants expressed concerns about data leakage, but their actions suggested otherwise. In both studies participants expressed a desire for greater transparency and control over information disclosure. Where none of this is particularly surprising, we were struck by the fact that sentiments expressed by survey participants about data collection and smartphone usage seemed to fall into three broad categories (see Table 1). We then explored more closely the sentiments expressed by both survey and interview study participants about data leakage and why they said they continued to use their phones the way they did. Our findings suggest that the notion of “privacy paradox” that is often used to explain these findings may obscure complex dynamics around technology use and data disclosure.

Tracking is understandable

Although the vast majority of responses in interviews and survey expressed significant distaste for tracking and third party data collection activities, some felt that there are legitimate reasons for data collection:

“I see the purpose of collecting data for improving the application behavior e.g. Amazon book suggestions, but I do not approve the selling of my data to third parties, especially when I am not getting a share of the revenue!” – survey, location undisclosed.

After all, people did want personalization of services, although they tended to identify these activities with expected, and thus contextually appropriate data usage. Others pointed to the typical business model of smartphone application and other online businesses, as they understood it:

“I understand that this information is a revenue stream for free apps - but I believe it has gone too far” – survey, USA

Theme	Code	N
Tracking is understandable	Collecting data is OK	1
	Developers need it	6
The concerned and the outraged	Should be illegal	8
	Users need to be informed / more control	7
Dejected acceptance and compliance	Nothing is free	6
	The user has to comply	14
Tracking is Disturbing/Creepy		50
Unconcerned "I do not care"		2

Table 1: Thematic coding of qualitative data from the survey (total N=89, some statements were coded in more than one category)

The reality of the information economy is such that in the course of mundane interactions with everyday technologies users move smoothly between consumption of services and production of content that is in turn monetized by service providers. This kind of information economy turns all data produced by consumers into commodities that are either fed back to the consumer through improved services and personalization or sold elsewhere. Many of our participants understood this dynamic quite well:

“I understand the appeal from the developer's perspective -- being able to collect large amounts of data about users could be crucial to targeted ads and later versions of an app.” – survey, USA

Alongside this understanding, most expressed concern for the amount of tracking and qualified their statements to

caution about the run-away train of the scale of data collection. Many participants said they didn't care if advertisers had data about their application usage or how many times they played a game on their phone during work hours. Yet a vague concern for how this data might eventually be used and interpreted remained:

"I'm ok with them collecting some information about me, but it is a slippery slope and the developer should only engage in the activity if they are willing to accept the consequences and responsibility of handling it properly." – survey, USA

Most of our participants called for greater transparency, regardless of whether they approved of data collection or not. Identifying what kind of data is collected, where it is delivered and how it might be used is no small task and most of the time it is impossible. Even the most technically savvy smartphone users often do not conceptualize quite the extent of data leakage.

The concerned and the outraged

In most cases, news of data leakage elicited expressions of mild discomfort and references to creepiness:

"I think it's creepy and makes me think about all the apps I have on my phone. I use flashlight very often and I find it disturbing that it can collect personal information off my phone" – survey, Iceland

In several cases, however participants expressed significant outrage. When asked why, in all cases they explained they had not expected data collection to be quite so rampant, but perhaps the level of outrage was associated with the kinds of potential outcomes they envisioned:

"This is completely ridiculous, I would not invite people into my closet, this is way out of line. No I don't find it appropriate to give up personal information in exchange for this game and that they don't need more approval than they apparently do." – interview, Denmark

If we conceive of smartphones as extensions of the self then the owners expect informed consent prior to entry into these personal territories. In the quote above, the interviewee equates the smartphone to a closed personal space (a closet) — entry into which would require high levels of intimacy and trust. People react to these breaches in different ways. Some uninstalled the offending application immediately, while others called for some sort of regulation:

"No such data collection and/or distribution should be permitted." – survey, Denmark

Curiously, only a few of our participants were able to articulate who it was that they expected to act as a gatekeeper for data collection and tracking. Several noted that this was the function of the app stores. Most, however, appealed to some sort of vague general policies or even laws that "ought to be in place." Many also called for greater transparency. After all, much of what goes on with user-

produced data, whether it is content on Google or data leaked by the Flashlight application on Android, is shrouded in mystery:

"There should be much stronger regulations against third-party data collection and people should be better informed about this risk of third-party data collection" – survey, Iceland.

Where the general sentiment here was that the practice of data collection and tracking should be somehow limited or stopped, only the very few actively limited their own app usage. Although many called for greater transparency, few admitted to reading the EULA before installing applications of their choice. In the vast majority of cases the same participants expressed another sentiment, that of compliance in the face of constant demands on data.

Dejected acceptance and compliance

Taking care of private information on smartphones takes effort; reading the EULA, scanning the phone and monitoring application updates which often involve increased numbers of permissions. Some of our participants stated that in order to use their smartphone they had to cave in and accept data collection. After all, what's the harm anyway?

"I don't know what can really be done [about data collection] because then you have to stop using the smartphone as a smartphone." – interview, Denmark

Others said that they had never experienced anything negative and therefore had lowered their threshold of privacy concerns. It has become a part of their expectations and the more that this continues to happen the more they will expect it.

"This is just how these things develop and this will only get worse, we can not change it. But I don't see how I can be personally affected by them knowing stuff about me." – interview, Denmark

Although people are willing to trade their personal data for a benefit their reaction is often not one of dismissal of privacy concerns but one of discomfort from the loss of control they feel. To some extent people want the option to be included in the negotiations, not to be robbed of their data.

"There is not much you can do about this. If you want the app you just have to accept this. Otherwise you are not using the phone the way it was intended." – interview, Iceland

Our participants continued with their habits after being exposed to the data leakage, saying that it is 'good enough'. Nine of the 13 interviewees told us that although the information they learned from us did not result in them changing how they used their phones, they didn't quite feel as comfortable about it anymore.

"The attitude has changed but the usage has not, I just feel a lot more uncomfortable using my phone" – interview, Denmark

The level of resignation was striking in most of these responses. More often than not, the explanation for why our participants didn't spend the time to read the EULA, overlooked that some of their favorite applications leak data and ignored the assortment of tools available for limiting data leakage of this kind was one of helpless compliance:

"It's unsettling and not ok, but I feel very powerless against it." – survey, Denmark.

Although we might expect people to exhibit high levels of concern for their privacy, how can we expect them to act on these concerns if they are convinced that there is not much they can really do? According to Altman [3] such lack of recourse can result in adjustments to the desired privacy level towards actual privacy level. People do so when they feel that they have no boundary control options to fend off repeated encroachments. After all maybe we are not looking at a privacy paradox. What we are seeing is a certain shift in norms and tensions between what our participants want, increased control and transparency and what they realize they can get. This is where our participants exhibit a pragmatic attitude and choose the benefits of owning and using a smartphone without limitations over the cost that comes with guarding their personal data from overprivileged applications. Smartphone users do not want to uninstall all of their overprivileged applications. They accept that this is payment for getting something for "free", and they feel at ease with paying for it as long as nothing negative happens, or as one respondent phrased it:

"I believe the adage: "When you think you're getting a free lunch; you're actually being served to someone else." It's a trade-off for free entertainment." – survey, USA

From privacy paradox to learned helplessness

Altman notes that repeated invasion into a primary territory can have serious consequences to a person's self-identity and inability to regulate access can in the long run cause a lack of self-esteem [3]. Invading a person's privacy is paramount to taking control of a person's life away from that person which can seriously affect an individual's autonomy and dignity [30]. According to Altman "it is a loss of control to others that is serious, not so much the mere exposure of information" [3]. Repeated invasions into a person's privacy and a conviction that there is no recourse can result in learned helplessness [23,30], when people stop responding to invasions even when presented with ways to defend themselves. Learned helplessness was originally identified in an experiment where a dog was put in an inescapable situation, as it was experiencing electric shock. Shortly thereafter, it was put into a different kennel where it could stop the shocks by performing a simple action. The dog did not attempt to evade the shocks; it just remained seated and stoically endured electric shocks, while the dogs in the control group quickly learned how to avoid the shocks [37].

The conception of learned helplessness was subsequently developed into a theory of helplessness and personal control

[1], applied to a variety of situations and identified as a significant factor in mental health. Learned helplessness typically happens when people come to believe that a situation is unchangeable or inescapable and will often construct reasons for why this is so even if solutions become available later on. Consider the following statements:

"I silently accept it. When you make me think about it, I kind of don't like it, but have probably forgotten all about it next time I download an app." – survey, Denmark

"It seems like a necessary evil at this point. Because it is so ubiquitous, I think that it's likely that this sort of thing will never go away." – survey, USA

In both cases there is an implicit agreement that the respondent has no way of affecting the situation and must accept it if they are to be able to go on. The implication here is that perhaps there are other explanations for what has been termed the privacy paradox [36], beyond decision-making conundrums and situational constraints.

DISCUSSION

When informed of information sharing practices of certain apps on their smartphone, our research participants for the most part expressed dismay, even outrage — but then proceeded with business as usual when it came to using their smartphones. To the directed advertising industry, this can be seen as good news, suggesting that consumers may protest information sharing but deep down do not really care enough to actually alter their behavior.

We take from our findings a different message: that at this stage in the smartphone era, many people are essentially creeped out (or, if they knew, would be creeped out) by the information sharing behaviors of the apps with which they have outfitted this extension of their personal space and primary territory. They may have formally agreed to it, and are able to rationalize their use of such apps when asked; and they may prefer to put the creepiness out of their minds in order to enjoy their moments of interacting with the apps; but this foundation of creepiness undermines and makes precarious the standing of the smartphone as part of daily life. From a business perspective, our finding is that any business model that depends on taking personal data from smartphone users without first establishing a solid basis for the level of emotional trust this entails is ripe for disruption.

We do not see the solution to these privacy issues in improving EULAs and in making informed consent more robust. Given that privacy concerns are generally rooted in complex social configurations, such approaches would seem to depend on people predicting that something might go badly and mitigating information flows on the front end. The reality is, however, that people have great difficulty predicting such outcomes because in fact there are far too many variables to consider - and the vast majority of times the negative consequences are the result of "not thinking" ahead. From a policy perspective then, the EULA hardly

qualifies as informed consent for this sort of data distribution to third parties. Instead, our analysis suggests that designers of apps and their data sharing policies need to confront the nature of creepiness head on. For this, we need a practical theory of creepiness, its varieties, and its temporalities (e.g., does creepiness fade over time with familiarity, and if so, what replaces it?). This paper is a first step in this direction but much more work needs to be done.

We also need apps that behave, and can be seen to behave, in ways that respect the user's personal space and the integrity of the context of use, much better than they currently do. Apps like High-Powered Flashlight or Fruit Ninja present themselves as simple, even trivial consumer products or services; people expect them to be simply fun ways of outfitting their personal space (in the form of the app space on their phone). Their creepiness lies in the realization that they *are more than they seem* – they are actually conduits for personal information to “leak” (intentionally) out of my personal space, indeed, out of myself. Upon discovering their underlying information gathering and dispersing nature, I suddenly find that I am actually sharing my personal space with unknown, and unruly entities that I did not really invite in, regardless of the EULA.

Telling people, in rational terms, that apps have this nature — of extending me beyond my personal space, and of allowing outside influences in — is likely not enough. Perhaps we (technologists) need to find ways to make people *feel* apps' active, connective nature in some visceral way. So far, the only *visceral* way people feel these sensations of discomfort and personal zone invasion/boundary collapse is what gets termed *creepy*. Are there more positive, visceral, affective responses that we could design for? Or can apps become effectively transparent in this regard — rather than a EULA, their design itself communicating a sense of access/thrill/risk (of both upside, and downside).

Altman posited negative psychological consequences if the person whose privacy is repeatedly violated is unable to regain control and to successfully manage access to the self [3]. From our data, we are not able to demonstrate that there are long-term negative consequences for consumers (e.g., encouraging lower commitments to personal privacy or placing creepy aspects of their personal environments out of sight and out of mind). But we suspect there are, perhaps through mechanisms that create and sustain bio-neurological stress, the harmful health consequences of which are increasingly being documented. We also acknowledge that creepy experiences may not always be negative and unwanted, as recent HCI research into “uncomfortable user experience” has begun to explore [6].

CONCLUSION

An analysis such as ours into negative emotional and moral reactions to smartphone app experience can be considered an exercise about values in design [13,20]. As our interview and survey respondents were able to articulate, their values

and the values implicit in the design and behavior of certain smartphone apps were not well aligned [21].

From a values in design perspective, what is at stake here is the kind of future we want to live in, beyond just how we wish people would use apps and what kind of responsibility we wish they could take for their own data. The issue of designing the kinds of sharing and disclosure applications is in part hitting the right balance between what can be seen and what can be kept secret, between control and automaticity. In part it is also about pushing changes in the norms around information sharing — we are fumbling with the now norms rather than looking to the where these norms might move as a result of technologies we are developing.

Cultural norms change over time, subject to a great many short- and long-term forces, technical, economic, social, and legal. Acceptable entertainment app behavior in 2013 will likely be different than in 2023. This will not eliminate creepy experiences, but will change the conditions under which they are encountered. Nor does the inevitability of cultural change diminish the need for responsible, value-sensitive design.

ACKNOWLEDGMENTS

This work was supported in part by the Intel Science and Technology Center for Social Computing. Big thanks to Jason Hong and his graduate students for providing support and advice in the initial stages of this work. We are indebted to all our respondents for their generous participation.

REFERENCES

1. Abramson, L., Seligman, M. & Teasdale, J. (1978) Learned helplessness in humans: critique and reformulation. *Journal of abnormal psychology*, 87, 1, 49-74.
2. Acquisti, A., & Grossklags, J. (2005). Privacy and Rationality in Individual Decision Making. *IEEE Security and Privacy*, 26-33.
3. Altman, I. (1975). *The Environment and Social Behavior*. Belmont, California, USA: Wadsworth Publishing company, Inc.
4. Barkhuus, L. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proceedings of CHI 2012* (Austin, Texas, 2012). ACM
5. Barkhuus, L., & Polichar, V. E. (2011). Empowerment through seamfulness: smart phones in everyday life. *Personal and Ubiquitous Computing*, 629-639.
6. Benford, S., Greenhalgh, C., Giannachi, G., Walker, B., Marshall, J., & Rodden, T. (2013). Uncomfortable user experience. *Communications of the ACM*, 56(9), 66-73.
7. Braunstein, A., Granka, L., & Staddon, J. (2011). Indirect content privacy surveys: measuring privacy without asking about it. *SOUPS '11*. NY, USA: ACM.
8. *Collins English dictionary*. HarperCollins Publishers, City, 2009.
9. Felt, A., Egelman, S., & Wagner, D. (2012). I've got 99 problems, but vibration ain't one: A survey of smartphone users' concerns. *SPPM '12* (pp. 33-44) NY: ACM.

10. Felt, A., Ha, E., Egelman, S., Haney, A., Chin, E., & Wagner, D. (2012). Android Permissions: User Attention, Comprehension & Behavior. *SOUPS '12 NY*: ACM
11. Flurry. (2013). Retrieved April 12, 2013, from Big Data: <http://www.flurry.com/big-data.html>
12. Fortunati, L. (2001). The mobile phone: an identity on the move. In *Personal & Ubiquitous Computing*: Springer-Verlag London.
13. Friedman, B. (1997) *Human values and the design of computer technology*. CSLI Publications, Stanford, CA.
14. Futuresight. (2011). Retrieved from User perspectives on mobile privacy: <http://goo.gl/9sgZXR>
15. Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., et al. (2005). Stopping spyware at the gate: a user study of privacy, notice and spyware. *SOUPS '05* (pp. 43-52). NY: ACM
16. Häkkinen, J. & Chatfield, C. 'It's like if you opened someone else's letter': user perceived privacy and social practices with SMS communication. *Proceedings of MobileHCI* (Salzburg, Austria, 2005). ACM
17. Karlson, A., Brush, A. J. & Schechter, S. Can I borrow your phone?: understanding concerns when sharing mobile phones. In *Proceedings of CHI 2009* (Boston, MA, 2009). ACM
18. Kelley, P., Consolvo, S., Cranor, L., Jung, J., Sadeh, N., & Wetherall, D. (2012). A conundrum of permissions: installing applications on an android smartphone. In *Proceedings of FCDS12* (pp. 68-79). Berlin: Springer-Verlag, Heidelberg
19. Kelley, P., Cranor, F., & Sadeh, N. (2013). Privacy as part of the app decision-making process. In *Proceedings of CHI 13* (pp. 3393-3402). NY: ACM.
20. Knobel, C. & Bowker, G. (2011) Values in design. *Communications of the ACM*, 54, 7, 26-28.
21. Le Dantec, C., Poole, E. & Wyche, S. Values as lived experience: evolving value sensitive design in support of value discovery. *Proceedings of CHI 2009* (Boston, MA, 2009). ACM
22. Lederer, S., Mankoff, J., & Dey, A. (2003). Who wants to know what when? Privacy preference determinants in ubiquitous computing. *CHI EA '03*. NY: ACM.
23. Lenhart, A., Ling, R., Campbell, S., & Purcell, K. Teens and Mobile Phones. *Pew Internet and American Life Project*. (2010).
24. Lin, D., & Loui, M. (1998). Taking the byte out of cookies: Privacy, consent and the Web. *Computers and Society*, 28 (1), 39-51.
25. Lin, J., Sadeh, N., Amini, S., Lindqvist, J., Hong, J., & Zhang, J. (2012). Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. *UbiComp '12* (pp. 501-510). NY: ACM.
26. Ling, R. (2008). *New tech, new ties: How mobile communication is reshaping social cohesion*. Cambridge, MA: MIT.
27. Maier, S., & Seligman, M. (1976). Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General*, 105(1), 3-46.
28. Muslukhov, I., Boshmaf, Y., Kuo, C., Lester, J., & Beznosov, K. (2012). Understanding users' requirements for data protection in smartphones. In *ICDE Workshops, 2012*. Washington, DC, USA.
29. Nafus, D. & Tracey, K. (2002) Mobile phone consumption & concepts of personhood. In J. Katz & M. Aakhus (Ed). *Perpetual contact: Mobile communication, private talk, public performance*. University Press, Cambridge
30. Nippert-Eng, C. E. (2010). *Islands of privacy*. Chicago ; London: The University of Chicago Press.
31. Nissenbaum, H. F. (2010) *Privacy in context: Technology, policy, and the integrity of social life*. Stanford Law Books, Stanford, CA.
32. Norberg, P., Horne, D. & Horne, D. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100-127.
33. Oksman, V. & Rautiainen, P. (2003) "Perhaps it is a Body Part": How the Mobile Phone Became an Organic Part of the Everyday Lives of Finnish Children and Teenagers. In J. E. Katz (Ed). *Machines that become us: The social context of personal communication technology*. Transaction Publishers, New Brunswick, N.J
34. Palen, L. & Dourish, P. Unpacking "privacy" for a networked world. In *Proceedings of the CHI 2003* (Ft. Lauderdale, FL, 2003). ACM
35. Peterson, C., Maier, S., & Seligman, M. (1993). *Learned helplessness: A theory for the age of personal control*. NY: Oxford University Press.
36. Rose, N. (1999) Powers of Freedom: Reframing Political Thought, NY: Cambridge University Press
37. Seligman, M. E. (1972) Learned Helplessness. *Annual Review of Medicine*, 23, 1, 407-412
38. Sundar, S., Kang, H., Wu, M., Go, E., & Zhang, B. (2013). Unlocking the privacy paradox: do cognitive heuristics hold the key? In *CHI '13 Extended Abstracts*, Paris, France: ACM
39. Vijayan, J. Flashlight app vendor settles with FTC over privacy violations. *Computerworld.com*, IDG Network, Dec 6, 2013, <http://goo.gl/OJ7P8P>
40. Xu, H., Luo, X., Carroll, J., & Rosson, M. (2011). The personalization privacy paradox: An exploratory study of decision-making process for location-aware marketing. *Decision Support Systems*, 51(1), 42-52
41. Zhou, Y., Zhang, X., Jiang, X., & Freeh, W. (2011). Taming information-stealing smartphone applications (on Android). *TRUST'11* (pp. 93-107), Heidelberg: Springer-Verlag.
42. Zafeiropoulou, A., Millard, D., Webber, C., & O'Hara, K. (2013). Unpicking the privacy paradox: can structuration theory help to explain location-based privacy decisions? *WebSci*, Paris, France.