The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI

Louise Barkhuus
Mobile Life @ Stockholm University
164 40 Kista, Sweden
barkhuus@mobilelifecentre.org

ABSTRACT
Privacy is a widely studied concept in relation to social computing and sensor-based technologies; scores of research papers have investigated people's 'privacy preferences' and apparent reluctance to share personal data. In this paper we explore how Ubicomp and HCI studies have approached the notion of privacy, often as a quantifiable concept. Leaning on several theoretical frameworks, but in particular Nissenbaum's notion of contextual integrity, we question the viability of obtaining universal answers in terms of people's 'general' privacy practices and apply elements of Nissenbaum's theory to our own data in order to illustrate its relevance. We then suggest restructuring inquiries into information sharing in studies of state-of-the-art technologies and analyze contextually grounded issues using a different, more specific vocabulary. Finally, we provide the first building blocks to such vocabulary.

Author Keywords
Privacy; user studies; location-based services; ubiquitous computing; online social networks

ACM Classification Keywords
H.1.2 [User/Machine Systems]: Human factors;

INTRODUCTION
With sensor-based and networked technologies built into the fabric of our everyday lives, from mobile personal devices to work places and public areas, personal information is now more accessible and sharable than ever before. This leads to a natural level of concern as to who has access to one's personal information and how it is shared. Research is therefore increasingly addressing issues of personal privacy. As opposed to data privacy (also referred to by Iachello and Hong as 'data protection' [21]), where laws and even constitutions prescribe what can and cannot legally be shared, personal privacy addresses the more fluid notion of privacy around a person, such as one's right to control personal information flow.

Issues of personal privacy in particular emerge in relation to the use of mobile devices that can take advantage of context information such as identity and location of a user, because mobile applications share information that was previously unsharable. Just as the digitization of music not only made it possible to edit music after recording but also enabled effortless (sometimes illegal) sharing, digitization of other information presents equally mixed advantages.

Many studies have investigated users' information management within interactive, mobile systems as well as individual's concern for information sharing when using mobile and sensor-based applications, studies that are relevant for CHI in terms of future design and development of technologies. However, despite research addressing overall concerns relating to personal privacy in these systems, there is a shortage of empirical studies focusing on the underlying contextually grounded reasons for people's privacy concern or lack thereof. Desiring to uncover simple rules and guidelines for design of sensor-based technology, many researchers have instead investigated narrow issues of information sharing such as *with whom* people would like to share particular information, and *when*, but few such studies have applied more theoretical notions of privacy to their data analysis or developed theoretical frameworks (with some notable exceptions [1, 12, 30]). With the limited analytical treatment of privacy, empirical findings are often locked into the context of the study. In this paper we argue that privacy is not an easily measurable unit and that we as HCI researchers and practitioners need to approach the notion through more contextually grounded measures. We hope to provide researchers with better insights into the concept of privacy and better tools to study this often slippery notion. We base our argument on Nissenbaum's theory of 'Contextual Integrity' [29], which views privacy as the appropriate flow of information rather than a static act of sharing. Illustrating key issues from the theory with our own empirical data, we provide evidence for the shortcomings of previous scenario-based studies of sensor-based services; we then provide recommendations for

CHI'12, May 5–10, 2012, Austin, Texas, USA.

research approaches to privacy and suggest the first important distinctions within a new privacy vocabulary.

TECHNOLOGY AND INDIVIDUALS' PRIVACY SENSITIVITIES

Researchers have already expressed concerns about digital technologies' ability to easily distribute even simple personal data such as birthplace and date; it has been argued that because electronic data can be shared more easily than paper-based directories, new standards for protecting personal information were needed [28]. In the mid-nineties, new telephone systems capable of displaying the caller's number (caller ID) were challenged as privacy invading, as they left the caller unable to remain anonymous until he/she chose to reveal him/herself [26]. More recently, sensor-based personal technologies such as location tracking applications and public surveillance have presented new challenges to the individual's sense of personal privacy. A number of studies have explored individual preferences for sharing information (such as location) through survey instruments, usually involving descriptions of likely scenarios of use [5, 25]. In Tang et al.'s study user-created data was used to generate scenarios about which participants answered questions regarding location-sharing preferences [35]. Using interview methods, other studies have explored issues such as attitude toward the sharing of online media [20], or WiFi snooping of content from personal laptops [24]. These studies set the scene for understanding initial reactions to information management, and provide a valuable foundation for later research. Yet, a basic challenge in this work is the inherent ambiguity in the definition of privacy itself, as well as questions over the generalizability of specific results.

Where most of these studies look at particular aspects of privacy concern, others have attempted to take a broader, more theoretical stance. Palen and Dourish, for example, provide a thorough analysis of privacy issues within socio-technical systems. Leaning on Altman's theory that describes privacy as dialectic and dynamic control processes for privacy boundaries, the authors analyze several relevant prototype and implemented systems. They point to the disparity that privacy regulation is not static or rule-based, yet technology, by definition, relies on both rules and stable functionality. The authors emphasize that the importance is not the specific privacy settings within the technology itself, but rather how the technology fits into cultural practices of privacy management [30]. Dourish and Anderson similarly argues that privacy cannot be viewed only as economic rationality, where personal information is provided in exchange for the benefit of a service or social value [16]. An interesting proposition following this argument is to view privacy as based on accountabilities of presence [37]. Our approach to privacy builds partly on this work in that we agree with the importance of fitting technology into already existing cultural practices. While these discussions focus on the issues governing interpersonal privacy management as a way to inform researchers and developers more broadly in terms of technology design, we propose a more contextually grounded approach to actual empirical research and a more dynamic use of the concept of privacy within HCI.

Another relevant theoretical approach is that of Boyle and Greenberg [13]; the researchers provide a taxonomy for different technical types of support for privacy in the use of video media awareness systems and describe a vocabulary for individual's modalities of privacy control. This vocabulary provides a useful and clear description of *how* individuals control their privacy in relation to media awareness. In this paper, however, we address the basic issue of *why* privacy is so important to use of sensor-based systems in order to continue answering how people negotiate privacy in everyday situations.

Location-Based Technologies

These issues of privacy are particularly relevant to location sharing. Location-based systems have entered many people's lives in the form of friend finders (such as foursquare, Gowalla and Facebook), map based searches (as provided with for example Google maps) and way-finding systems (GPS navigators). Such applications introduce a new set of privacy issues related not only to surveillance but also to simple factors such as control of impression management and interpersonal privacy [30]. It is therefore no surprise that many studies of people's privacy concerns and preferences for data sharing have focused particularly on location. Where some studies have based their findings on scenario testing [5], others have deployed systems with real-time sharing of location [36] or focused on commercial systems [20, 27]. These efforts to pinpoint aspects of people's concerns regarding location sharing are of particular interest in this paper. Where, from a naïve perspective, a person's location is not sensitive information (in that a person is generally physically visible to others), the recording and sharing of this information with people outside the physical vicinity make this information potentially sensitive. Researchers have, since early prototype systems, attempted to determine and predict specific location sharing practices [14, 25] and attributed results to the concept of 'privacy preferences'. However, most technology-focused studies of personal privacy preferences/perceptions have yet to provide a clear description of how these findings fit into a broader understanding of personal privacy.

The Understanding of Privacy

Nissenbaum's theory of *contextual integrity* [29] emphasizes that there is no aspect of human life that is "not governed by [context-specific] norms of information flow" [29, p. 119], those being cultural, ethical or moral norms. We are able to fluently transform our behavior according to these norms just as we appropriate our actions with people we have different relationships with. She argues that there are no such thing as universal privacy norms but that these are distinct to each situation, and assist in maintaining

contextual integrity. *Contextual integrity* describes a desirable state that people strive towards by keeping perceived-private information private according to the context. For example, people expect to share medical information with doctors but not with employers. Where it in some cultures yearly salary is perceived as private, within others it is normal to share this information. Contextual integrity thereby explains how privacy is grounded in each context, governed by pre-existing norms and values. To explain her theory of contextual integrity, she highlights three principles that dominate the public discussion of privacy policy: (1) protecting the privacy of individuals against government intrusion, (2) restricting access to sensitive personal information and (3) protecting personal space. It is argued that because many new socio-technical and sensing technologies (for example public surveillance) fall into grey areas or even outside these principles, there is no pre-defined understanding of privacy in these situations. The argument relates to Boyle and Greenberg's explanation of "inadvertent privacy violations" where the media space design aligns poorly with social and human factors [13]. It thus appears rather narrow to attempt to generate generalized, rule-based principles about personal privacy preferences. Understanding personal privacy concern requires a contextually grounded awareness of the situation and culture, not merely a known set of characteristics of the context.

The current popular attempt to provide general guidelines about privacy preferences is flawed exactly because of the user studies' lack of situational foundation. We suggest that other, more theoretical, measures are included in the broader discussion of privacy in relation to technology studies. In the following section we expand on Nissenbaum in relation to our own empirical data to support our argument.

THE CASE FOR CONTEXTUAL INTEGRITY

To support our proposition that privacy should be approached as a much more flexible notion and to explain the theory of contextual integrity, we lean on empirical data focusing on mobile personal information management. The data is from a project studying university students' use of a mobile social network (Facebook); this group is of particular interest since they illustrate common use of a network that is highly integrated into their daily lives. We do not claim that these students represent the broader population but rather that this sub-population represents a particular set of values in an extreme, highly social, highly mobile setting. Their practices illuminate aspects of privacy that so far have been ignored. The project itself focuses on broader issues; here we focus on their information-sharing and privacy perceptions (see [12] and [6] for further details on the project). The data consist of qualitative semi-structured interviews with 60 students as well as data collection from their online profiles. All participants accessed their online social network mostly from their mobile device (iPhone, Blackberry), resulting in very integrated use where other types of text communication such as text messaging were placed side by side in terms of importance.

In this section we examine three parts from Nissenbaum's theory in an aim to provide real-life examples of how privacy is perceived, negotiated and articulated: social appropriateness, distribution and change of norms. We apply our own data selectively, acknowledging (and emphasizing) that one particular online social network does not contain generalizable behavior,. We aim to exemplify particular characteristics of Nissenbaum's theory to build our argument.

Social appropriateness

Part of Nissenbaum's reasoning refers to how norms of appropriateness determine what personal information is fitting to reveal in a particular context. Each social setting prescribes what kind of information is expected to be shared; it is appropriate to share medical data with doctors but not with distant acquaintances and it is normal to share romantic details with friends but not with colleagues. These norms are primary to behavior and inform the contexts wherein actions are performed. Among our participants it was clear that such concept of social appropriateness was in place regarding their use of the online social network. Just as face-to-face contexts are governed by norms, so were the online communication situations that the participants reported. This communally understood behavior, or as Mancini et al. refer to as 'unspoken code of conduct' [27], contributes to a general understanding of social context, and participants had few problems with others overstepping their bounds of privacy. In the few cases where such overstepping had occurred, swift action by the participants had been taken, such as unfriending a friend who repeatedly sent unwanted messages. This general understanding of limits and bounds was in place partly because the network had been part of the participants' lives for 3+ years but more importantly because they had the ability to weigh situations and act accordingly.

Nissenbaum also stresses how an important part of norms of appropriateness is the distinction between different relationships, different social roles. How a person behaves is dependent on the other people present in the context. For our participants it was clear that their personal information sharing was grounded in their own perception of their social relations present on Facebook. Their content sharing was meant for their full set of friends with few exceptions. As in other research [20], none of the participants had divided their friends into sub-groups, providing reasons that this is a time-consuming task or due to unawareness of this feature. Instead they were reasonably aware who was on their friend list and shared information accordingly. The 'mom test' was not uncommon, illustrated by this male participant: "I am friends with my mom on Facebook so I guess I am okay with my mom seeing it." Participants' content was self-censored to fit the 'greatest common divisor'.

Another method of tailoring the content to the socially appropriate situation was to publish content that was specifically understood by a subgroup but not by the rest of the group. While the study took place the university held an 'end of semester' extended weekend and although it was known internally by students that a lot of drinking took place, the status messages were cryptic enough to leave outsiders with some ambiguity as to what was going on. Participants wrote content such as: "[name of special weekend]…Worlds greatest college weekend…love you" and another participant wrote: "thinks [name of special weekend] is going great =]". The plausible deniability that vague messages provided was enough for the participants to find them acceptable, even though the vast majority stated that alcohol in pictures or status message was improper, especially because two thirds were under 21, the legal drinking age in the US. It was necessary for the participants to manage their content tightly due to their broad set of friends, conforming to the notion of social appropriateness.

Challenges to Research

Nissenbaum points out that although norms of appropriateness are part of everyday life, these norms are not explicitly addressed in research that informs privacy policies. We also find that much research relating to privacy issues and technology has, if not simply ignored context, then attempted to generalize to other contexts on the basis of characteristics. It is not uncommon for studies to inquire to their participants' preferences for information sharing on the basis of scenarios or personas rather than within an actual sharing situation. Before commercial sharing applications were available this seemed acceptable, but even with experience sampling made possible in networked mobile devices, the 'experience' tracked is still too complex to generate general rules from. Our concern here is not simply the quality of research methods but the approach to privacy that many studies take, attempting to gather information that can inform design in the form of 'people's general preferences for privacy'. By studying behavior as predictable and generalizable, personal privacy research is missing out on opportunities to understand people's underlying motivations for sharing and not sharing in the first place. This brings us to a discussion of the distribution of information, or information flow.

Information Flow

Another set of norms that Nissenbaum uses to explain privacy in terms of contextual integrity is the concept of information flow. Like the norms of social appropriateness, existing norms of information flow are engrained into everyday life, again in particular related to different relationships. Relationships are in essence defined by what type of information we share with one another. People share more intimate information with close friends and more general information with acquaintances. What we tell people about ourselves is in a sense as important as what we hold back in terms of forming relationships, a factor that should inform our understanding of privacy. These norms of information flow help explain the sudden concern over sharing personal information such as present location or preferences for particular material commodities. Yet, most people also have an exhibitionist desire to share at least some personal information with different sets of people, described as "self-exposure".

Reasons to share personal information

Before the advent of online social services such as Flickr, Twitter and Facebook, many people found it difficult to imagine the value of sharing with superficial relations trivial, personal information as is done now through status messages and photos of small everyday events (and for many it still is). But interest in others' personal lives have always been part of human nature, and with the introduction of reality shows in the late '90s more and more people have been getting their '15 minutes of fame', to the joy and entertainment of others. What appealed to our participants were some of the same factors that are at play in reality television where apparently normal people (as opposed to people with particular abilities leading to their fame such as singing, acting, writing or sports) have their personal lives exposed. It is the dual pleasure of both being able to expose ones' self and being able to 'snoop' on others' private lives that is at play here. A female participant for example expressed why she found Facebook so intriguing: "It's for the same reason that, you know, people find celebrities entertaining and so many gossip, you know, it's like reading about other people's life, especially when you know them, is just really entertaining". Exposing personal information in essence makes the person 'feel famous', especially when others comment on content (in public or semi-public), which is re-affirmed as the person is able to see others' personal information and comment on it. The network, in this sense, fulfilled a need for self-exposure but at the same time users were able to manage their personal information. The dual desires to expose private personal information to seek attention and to 'snoop' on others' personal lives work well in an online social network where relationships are dominated by weak ties [17].

Reasons not to share personal information

Obviously there is also personal information that people are not interested in sharing with (selected) others and this is the area of which most research has addressed. Many studies, particularly of ubicomp technologies, focus on what information people do not wish to share. However, what many of these studies ignore and misunderstand is the distinction between different intents behind not wanting to share a piece of personal information. A participant who responds that they would not share a particular piece of information with a particular other is not necessarily *worried* about the other person obtaining this information or expressing privacy concern. There are multiple contextual factors at play. Here we propose a simplified explanation where there are two reasons for not sharing information, one a genuine concern about others knowing something (a *secret*) and the other simple *modesty* or politeness, not

wanting to 'burden' others with seemingly insignificant information. Where *modesty* governs particularly weak ties (e.g. we do not have genuine *concern* about the person next to us in the coffee shop knowing our name and position, yet we would not walk over and volunteer that info); *secrets*[1] govern to a higher degree close ties (e.g. we would not expect our close friend to tell our spouse what we have bought them as a Christmas present). The closer a relationship gets, the more information is often exchanged between the partners and the more potential 'secret' or 'relationship specific' information exist. In an odd way new social media technology supports the *modesty*-motivated desire not to share, by making it available for easy access rather than having to specifically provide the information. While people are not sharing information specifically with each other, they are making personal information available to interested others; it is not a motivation to hide information but a motivation to adhere to common practices of what information to (explicitly) share with specific people in specific situations (information flow based on relationships). It was clear from our study that the participants were able to negotiate their shared content to fit with both motivations of modesty (by not 'forcing' information on anyone) and secrets (by not revealing inappropriate information to people on their friend list).

Negotiating information sharing on the ground

As we have argued, supported by Nissenbaum, information flow defines individual social relationships, but digital technologies are not very good at managing groups of relationships, except laboriously. Facebook does have an option to separate friends into groups and publishing content to specific groups only; however, despite its apparent usefulness, none of our participants had taken advantage of this feature, either because they were not aware of it or because they did not want to spend the time on customization. Perhaps, one of the problems with this tool is that a friend list is still static when it has been created. Following the principles of contextual integrity, each status message and each picture should be published to a new unique list of friends. Another obvious flaw, based in the user interface, is that receivers of the messages cannot see who the message had been shared with, which subset of friends or the sender's whole network; this prevents the message from being socially translucent [18]. Instead of using friend subsets, the participant showed us how they negotiated the sharing by going by 'greatest common devisor. When asked if they worried about their privacy, a male participant said that he did not but "I am conscious of what I put up there", indicating that he would rather censor the content than fiddle with detailed privacy settings. This highlights the point that controlling information flow is negotiated on the ground, locally, instead of through a complex set of settings behind several pages, which was not transparent to view at each posting[2].

The area where most self-presentation management took place for the participants was in relation to pictures posted on Facebook. One female participant for example said: "People that take pictures of me are people who are allowed to take pictures of me I guess, like friends..." and another participant said "If there is a picture where I don't look my best I don't want to be tagged and I will untag it." The same participant expressed how she was able to check out her boyfriend and how she would look through his profile on occasion: "…there have been a few incidents were, like he will be tagged in a picture with another girl and I will ask 'what was that?' and it is nothing that has been too problematic…" Identity management was conducted as an ongoing activity where the personal information sharing was constantly held up against the social norms of information flow, norms tailored to the participants' relationships.

Research often points to the discrepancy between expressed privacy concern and apparent oversharing, or inflated self-reports of concern that do not fit with actual information protection behavior [1]. We found part of a potential explanation in that several participants were not sure exactly who they where sharing their content with. One participant, for example, replied to our question of which friends he shared content with: "I think… I don't even really know… I think it is just private, I mean only people from [the university] can search me or my other networks…" It was not uncommon to be unsure of their actual settings, which in return questions, if not the validity, at least the usefulness of studies viewing 'friends-only' sharing as a 'privacy enhancing behavior' [34]. Instead, we note how information-sharing practices were negotiated on the ground, within single experiences; if our participants experienced something negative, they changed behaviors.

This data should clearly illustrate our earlier emphasis that studying privacy preferences is complex and should be rooted in context. No matter how detailed a scenario is given within a study and no matter how many characteristics of participants' real life is based on real information, each situation is unique for the individual, and each decision to share or not is decided within that context.

Change of Norms

As a final characteristic from Nissenbaum's theory, we turn our attention to the changing of norms over time. What was appropriate to know about other people at one time is not the same today. Where in the 1990s it was inappropriate or at least questionable to be aware of the caller's identity

[1] By secrets we do not mean juicy tales of cheating spouses but simple daily relationship negotiations, such as not wanting a spouse to hear a heart-to-heart talk with a friend.

[2] Facebook has recently changed the way users chose who to share content with and selecting people and lists appears to be slightly more transparent.

before picking up the phone, it is today inconceivable not to have the caller ID function on cell phones. Especially with technological development, norms change rapidly, some more subtly than others.

Several of our participants made reference to changes of this sort. They had encountered something undesired, such as an unwanted post on their profile, and had de-friended the offender; their friend-accepting behavior consequently became more conservative, and most participants claimed that they knew virtually all of their friends personally and had met them in real life. We propose that even within an online social network norms change, and participants illustrated a change in social norms in their transition from high school to university. Where it was appropriate to accept (and possibly interact) with all friend requests in high school, now more moderate norms reigned.

One problem with inquiries into privacy preferences and concern in relation to prototype technologies is these technologies' lack of establishment among the people testing them. Much of the research within ubicomp and HCI uses prototypes, which can be a useful tool for receiving early feedback to new technologies, but which can be more complicated for studies predicting behavior manifested within repeated and socially established use. Looking at early preferences for potential sharing of certain data does not generate invalid findings but because their lack of grounding within actual lived situations they merely provide a tiny slice of potential human behavior in potential situations. It is therefore not only difficult to describe generic preferences among people but it is questionable how useful such prescriptions are long-term. At best they represent a snapshot in time, describing how current norms influence behavior.

We now step outside Nissenbaum's theory and continue by looking at a well-studied type of personal data in terms of concern for sharing: location. Ubicomp research, in particular, has maintained location as a focus of research, partly due to recent technological developments and partly due to its wide appeal in social applications. By illustrating the use of location within our own data, this discussion adds to the argument that preferences for personal privacy are grounded in contextual integrity.

LOCATION SHARING AND PRIVACY

Location has recently become a type of personal information that is possible to record and share, in real time, through our pervasive mobile devices. This intriguing type of personal information can have useful and productive applications (such as locating emergency calls from cell phones, locating nearest point of interest and locating friends of interest), but can also aid misuse (locating victims by criminals and stalkers, targeted but unwanted advertisements, and locating by unwanted friends/family). More importantly, the very essence of personal privacy includes the permission to 'hide', that is, placing oneself in seclusion, away from self-selected people. The prospect of being located is therefore also distressing for people if not in full control of this information. Referring back to Nissenbaum's concept of contextual integrity, it should be clear how general societal norms, in terms of the appropriateness of real-time location sharing, have yet to develop; so far we adhere to common sense. It is therefore no surprise that research has been particularly curious about people's concern for location sharing.

Many studies of privacy concern relating to location have been scenario based; participants were provided with descriptions of situations (often based on person-relevant data) and then asked if they were comfortable sharing particular location information with particular people [5, 10, 11, 25, 35]. Others have employed actual sharing technologies and in-situ studies [4, 22, 23, 31, 36, 38]. Each study focuses on particular elements of behavior or perception, for example how *much* concern individuals express [5], how their level of concern influences their privacy settings [36, 38] or people's willingness to share location information with others depending on location and recipients [2, 14, 25]. Most of this work attempts to pin down people's perceptions and preferences of sharing personal data with a wide, but specified, set of people.

Different Types of Location

In these studies the concept of location is addressed as *a single type* of information. Papers conclude and generalize about 'location privacy' based on the location type that they are looking at, despite other studies and other applications using and studying different granularities of location as well as different types of location information. Where some applications rely on exact geo-referenced locations using latitude and longitude, others use user-entered place names. Yet others (for example commercial 'check-in' services such as Gowalla, foursquare and Facebook's check in) use community-named locations. The level of accuracy also differs, but with most commercial systems now using GPS, with WiFi fingerprinting as backup, the imprecision of early systems has decreased. Unfortunately, treating such diverse information as a single entity waters down analysis and generalizability between studies. Systems and findings are compared inappropriately to other systems and research, leading to analogies and conclusions that may not be justified. 'Location' is, like the notion of privacy, a vague concept, interpreted variously according to context.

One way this problem manifests is that reviews of literature on location sharing services often do not actually mention what type of location data was shared. The difference between sharing a named location (such as 'Home', 'Peet's Coffee', 'University and 9th Av') and exact geo-referenced location is significant for social appropriateness; where providing a self-defined location is better in upholding accountability and contextual integrity, a geographical reference is more likely to reveal locations with no social translucence [18]. Similarly, the difference between 'checking in' and constant background updating also poses

a significant shift in accountability; a manual check-in is controlled by the user, but a location provided by background updating can be misunderstood by even the closest friends. This is particularly relevant when adhering to the view that privacy is based on accountabilities of presence [33, 37].

Some studies do address the issue of disclosing location at different levels of granularity; Consolvo et al., [14] for example, directly asked people about which level of location information they would reveal and Hecht et al., looked at manually entered location [19]. There is a major difference of the user's perception between sharing location with different granularity and of different type, and little research so far has addressed this issue. Comparing or contrasting people's preferences for sharing geo-referenced location with sharing self-defined, text-based locations makes little sense, owing to social norms around providing this information [32]. In fact, the increased social translucence of self-defined location not only adds contextual information, it can also contribute to maintaining personal relationships [4] and is therefore likely more rewarding to use.

Motivations for Sharing Location

Although some studies have looked into motivations for sharing location digitally with friends or sharing with third parties [6, 14, 15], few studies have been able to ground their studies in real situations. In the studies that have investigated why and how their participants wanted to share location information in particular situations, the situations were mainly provided as scenarios, or location was given only as a geo-referenced type [5, 11, 25, 36].

Our own data confirmed that individuals occasionally use Facebook to post their location. Yet, their understanding of social appropriateness prevented them from just showing up at a friend's location unless explicitly invited. "That would be creepy," participants explained. The norms in our (Western) society (at least since the telephone was widely adopted) dictate that people do not 'just show up' but make prior social engagements. This transferred to Facebook even though the knowledge of people's whereabouts was now often more detailed. Interestingly, this use is in contrast to social networks that were specifically developed for meet-ups such as Dodgeball[3]. We have earlier pointed out how specifically inviting status messages can lead to real-life meet-ups, but that these were almost always mediated by other communication [6].

This common rule in turn made location a fairly uninteresting status to post for our participants unless it was in the context of something more 'status-worthy' and we found less than 5% of their status messages contained reference to present location. The surprising finding that none of the participants 'checked in' (through Facebook's own feature) or pushed Gowalla or foursquare locations to their Facebook profile was very likely a feature of our student sample. They lived in a small college town with few core social spots and possibly felt the need to convey their identity [15] through other expressions.

Inquires into people's 'privacy preferences' and 'privacy practices' are often based on the belief that privacy is of clear concern to all. But while people of course do have such concerns, these are rarely well articulated or generalizable, even within individuals. Studies often inquire what types of information the participants are willing to share, and in what situations, but the analyses often neglect to inquire the deeper reasons why people are not willing to share particular information. Studies that have asked this question often find different reasons than those of privacy. Consolvo et al., for example, in their initial study of location-sharing preferences, highlights issues of privacy between hierarchical relations; employees are more cautious about sharing location information with their managers, which would then indicate a generic predictor with whom people want to share information with [14]. However, in the authors' later technology implementation of Reno, they point out that participants mainly turned off automated sharing in order not to 'spam' the people they shared location with [22]. Bales et al. also found that the 'off setting' in their location-sharing system CoupleVibe was used merely out of courtesy for the sharing partner, for example due to a known time difference, instead of a clear preference for privacy [3]. And in our own study of Connecto, groups of co-workers, including a manager, shared location throughout the study without expressing any significant concern [4]. This ties back to the discussion of modesty versus secrecy in terms of motivations. Even these few examples show how modesty is also too simple a concept to explain motivations; so where people might seem reluctant to share location data by answering negatively in a study or survey, the reasons for this are contextually defined. Unfortunately the data from the many studies that look at sharing preferences skews our knowledge of real users' motives and more broad considerations for their privacy.

We now turn to a discussion of possible implications for research that this rather rigid notion of privacy in HCI research is contributing to.

IMPLICATIONS FOR RESEARCH

As described in this paper, the actual data that are often shared, or imagined shared, are in no case homogenous across studies; where some share a community-defined location, others share an exact geo-referenced location, etc. This is in itself not problematic because different systems make use of different types of data, but when researchers compare findings across studies, the fallacy is to think that these can inform one another and together build a

[3] Dodgeball was a text-message based social network existing from 2000 to 2009. It enabled people to see which of their friends were nearby, focusing on social gathering.

generalized notion regarding 'privacy behavior'. It is for example impossible to ask if people have any concerns about sharing their 'location' and get a useful answer, since location is not a discrete piece of data; it is an interpretive type of information that can be perceived differently by different people depending on a range of details; it is always contextually dependent. A major gap in the discussion is therefore one of data type. People's answer to 'where are you' depends on the situation, as witnessed by the location-sharing studies that were based on self-input locations (such as [4] and [22]). Sometimes the importance is a specific restaurant, other times it is simply the city. The choice depends on numerous factors, all grounded in the specific context of the information provider and his/her receivers. Unfortunately, geo-referenced systems are not flexible enough to support different granularity, apart from simple obfuscation. Geo-referenced systems therefore do not provide well for individual accountability; in a world where the pub is next door to the library, the goal to uphold contextual integrity can easily fail. We therefore suggest that location-sharing studies emphasize granularity and data type of location before trying to generalize uses and sharing preferences; by examining location sharing in context one can better understand the diverse set of privacy issues.

Our second implication can be found in our finding that many studies are not grounded in actual user situations. Although many studies use real user data, it is rare to see studies that implemented real systems with real data sharing or which used in-situ survey data. Instead users are asked to evaluate potential situations on the basis of their own data. Although these studies might introduce us to basic issues of personal information management and hint at early reactions to sharing potentially sensitive data, they should not be mistaken for rule-generating or theory-informing research. When, as Nissenbaum suggests, all decisions and potential privacy violations are based in the real-life context informed by social norms, it is not possible to generalize further than the studies' own context. We therefore call for more contextually-grounded research that explores privacy issues in the wild.

Third and finally, the issues we as researchers should address regarding privacy are more complex than just *concern*. Privacy is also a way for us to maintain distance to weak-tie acquaintances and it is something we maintain by not 'bothering' others, as illustrated above. There are many reasons for hesitancy to share personal information, from courtesy and modesty to uncertainty about the audience. Few studies within HCI and ubicomp today have been able to pinpoint exactly *why* their participants reacted the way they did, only answering *how* they reacted and *what* their practices were. If we view privacy in a broader light, it transforms (as Palen and Dourish has also pointed out [30]) into the everyday negotiations done to manage not only the availability of one's personal information, but also the information that one receives about others. This important characteristic of privacy leads us to emphasize that it is the *appropriation of behavior* in the situation that informs new research and not a behavioristic belief that people's actions are based on a structured set of privacy concerns. Instead of focusing on the *how* and *what* in terms of people's preferences for personal data sharing, we need a foundation of research that looks at *why*.

CONCLUSION

In this paper we have attempted to frame and discuss the notion of privacy as it is studied and presented in relation to HCI and ubicomp technologies. It is no doubt a slippery notion that tends to be used for any kind of concern about personal information recording and sharing. With location detection being of particular interest, many studies have examined concern for the sharing of personal location data, real-time as well as recorded. We believe that a great portion of this research is not grounded in real-world situations and is therefore unable to provide more than coarse generalizations, which in many cases are easily predicted by applying already-present theoretical frameworks such as Nissenbaum's contextual integrity [29], Altman's theoretical privacy framework (as used by Palen and Dourish [30]) and Agre's discussion of 'the new landscape' [1].

With illustrations from our own study of practices among highly mobile users of an online social network, we emphasized how articulated privacy concerns were less visible among our participants but that they instead adjusted their practices according to social norms. They changed behavior (e.g. being more conservative in their friend acceptances) mainly after having experienced something negative, something that made them question their private integrity. Although we do not claim that our data set, collected from a very homogeneous and non-representative set of people, provides generalizable findings about broader aspects of privacy, we found it useful to illustrate some of the points of our argument.

From our analysis of literature and our own empirical data it is clear that privacy concerns are not easily measurable. People want accountability for their actions, they want control over their own personal data, and they want plausible deniability for their recorded and shared personal data. When data is already semi-public, for example when a person is out in public or self-posts information online, sharing does not violate any sense of privacy, except if this sharing goes beyond what is commonly conceivable in accordance with contextual integrity. It is therefore reasonable to state that people prefer manual systems where they keep control over their data and in which they can define the shared data themselves, from situation to situation (such as for example self-defined labels). At the same time it is important not to underestimate the social consideration people have for one another, particularly in close relationships (strong ties); they are simply not interested in providing more information than necessary in scenarios of personal information sharing due to modesty

but also because this can also lead to complex scenarios of sense of reciprocity (see Bales et al., for a discussion [3]).

Our major aim with this paper is to call for a more nuanced treatment of the notion of privacy within HCI. In fact we will go as far as to suggest minimizing the use of this loaded and commonly-misunderstood word. Often, research is in fact addressing issues of information sharing practices and concern for revealing personal data. Using the word *privacy*, particularly to elicit data from participants, not only introduces the complexity of multiple interpretations of the word but also automatically stamps the data as sensitive. Location information, for example, is in many public situations the most available data about a person, free for all to see; yet many current studies present a general concern for sharing such information widely.

We therefore suggest that future studies use a different, more specific vocabulary for empirical research. Such vocabulary should explicitly distinguish between carefully determined reasons *why* individuals want to share or not share information before looking at how and what. Our initial analysis here uncovered the notions of secrecy and modesty, concepts that provide a first step for such distinction. We suggest that the next steps are taken in the direction of uncovering the established sense of social appropriateness and its influence on people's information sharing practices. Yet such investigation should distinguish between social norms and personal appropriateness due to their distinct characteristics and influence of behavior. Viewing social norms and personal appropriateness as fluid notions will prevent overarching generalizations based on situational characteristics and instead provide a better understanding of why, rather than how, individuals want to share personal information. Only by utilizing such detailed vocabulary can we begin to gain a better insight into personal information sharing practices.

ACKNOWLEDGMENTS

The author would like to thank Nis Bornoe for help with the fieldwork. Also thanks to Leysia Palen and Paul Dourish for thoughtful comments on the paper. Finally, thanks to Valerie Polichar for editorial assistance.

REFERENCES

1. Agre, P. Introduction. In Agre, P. and Rotenberg, M. *Technology and privacy: the new landscape.* MIT Press. 1998.
2. Anthony, D. Henderson, T. and Kotz, D. 2007. Privacy in Location-Aware Computing Environments. IEEE Pervasive Computing 6, 4 (October 2007), 64-72.
3. Bales, E., Li, K. A. and Griswold. W. 2011. CoupleVIBE: mobile implicit communication to improve awar eness for (long-distance) couples. In Proc. of CSCW '11. ACM, New York, NY, USA, 65-74.
4. Barkhuus, L., Brown, B., Bell, M., Sherwood, S., Hall, M. and Chalmers, M. 2008. From awareness to repartee: sharing location within social groups. In Proc. CHI '08. ACM, New York, NY, USA, 497-506.
5. Barkhuus, L. and Dey, A.K. Location-Based Services for Mobile Telephony: a study of users' privacy concerns. In Proceedings of Interact 2003, pp. 207-212.
6. Barkhuus, L. and Tashiro, J. 2010. Student socialization in the age of facebook. In Proc. of CHI '10. ACM, New York, NY, USA, 133-142.
7. Bell, G. and Dourish, P. 2007. Yesterday's Tomorrows: Notes on ubiquitous computing's dominant vision. *Personal Ubiquitous Computing.* 11, 2, 133-143.
8. Bell, M., Reeves, S., Brown, B., Sherwood, S., MacMillan, D., Ferguson, J. and Chalmers, M. 2009. EyeSpy: supporting navigation through play. In Proc. of CHI '09. ACM, New York, NY, USA, 123-132.
9. Bell, M., Chalmers, M., Barkhuus, L., Hall, M., Sherwood, S., Tennent, P., Brown, B., Rowland, D., Benford, S., Capra, M. and Hampshire, A. 2006. Interweaving mobile games with everyday life. In Proc. of CHI '06. ACM, New York, NY, USA, 417-426.
10. Benisch, M., Kelley, P.G. Sadeh, N. and Cranor, L. F. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. In *Personal and Ubiquitous Computing*, December 6., 2010, 1-16.
11. Bernheim Brush, A.J., Krumm, J. and Scott. J. 2010. Exploring end user preferences for location obfuscation, location-based services, and the value of location. In Proc. of Ubicomp '10. ACM, 95-104.
12. Bornoe, N. and Barkhuus, L. 2011. Online Social Networks On-The-Go: An Exploration of Facebook on the Mobile Phone. Horizon paper in Proc. of CSCW EA 2011.
13. Boyle, M. and Greenberg, S. 2005. The language of privacy: Learning from video media space analysis and design. *ACM Transactions of Computer-Human Interaction.* 12, 2 (June 2005), 328-370.
14. Consolvo, S., Smith, I. E. Matthews, T., LaMarca, A., Tabert, J. and Powledge, P. 2005. Location disclosure to social relations: why, when, & what people want to share. In Proc. of CHI '05. ACM, 81-90.
15. Cramer, H., Rost, M. and Holmkvist, L.E. 2011. Performing a check-in: emerging practices, norms and 'conflicts' in location-sharing using foursquare. In Proc. of MobileHCI '11. ACM, 57-66.
16. Dourish, P. and Anderson, K. 2006. Collective Information Practice: Exploring privacy and security as social and cultural phenomena. Human-Computer Interaction. 21, 3 (September 2006), 319-342.
17. Ellison, N. B., Steinfield, C. and Lampe, C. 2007. The Benefits of Facebook ''Friends:'' Social Capital and College Students' Use of Online Social Network Sites. Journal of Computer-Mediated Communication 12 (2007), pp. 1143-1168.

18. Erickson, T., Smith, D. N., Kellogg, W. A., Laff, M., Richards, J.T. and Bradner, E. 1999. Socially translucent systems: social proxies, persistent conversation, and the design of "babble". In Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit (CHI '99).

19. Hecht, B. Hong, L., Suh, B., Chi, E.H. Tweets from Justin Bieber's Heart: The Dynamics of the "Location" Field in User Profiles. In Proc. of CHI '11. ACM, New York, NY, USA.

20. Hull, G., Lipford, H. R. and Latulipe, C. Contextual gaps: Privacy Issues on Facebook. In Ethics of Information Technology, 2010, Springer.

21. Iachello, G. and Hong, J. End-user privacy in human-computer interaction. Foundational trends of human-computer interaction, 1, 1 (jan, 2007), pp. 1-137.

22. Iachello, G., Smith, I., Consolvo, S., Abowd, G. D., Hughes, J., Howard, J., Potter, F., Scott, J., Sohn, T. and Hightower, J. et al. Control, Deception, and Communication: Evaluating the Deployment of a Location-Enhanced Messaging Service. In Proc. of UbiComp '05, Springer, 213-231.

23. Jedrzejczyk, L. Price, B. A., Bandara, A.K. and Nuseibeh, B. 2010. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In Proc. of SOUPS '10. ACM, New York, NY, USA, Article 14.

24. Klasnja, P., Consolvo, S., Jung, J., Greenstein, B.M., LeGrand, L., Powledge, P. and Wetherall, D. 2009. "When I am on Wi-Fi, I am fearless": privacy concerns and practices in everyday Wi-Fi use. In Proc. of CHI '09. ACM, New York, NY, USA, 1993-2002.

25. Lederer, S., Mankoff, J. and Dey, A.K. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. In Proc. of CHI EA '03. ACM, New York, NY, USA, 724-725.

26. Lyon, D and Zureik, E. Surveillance, Privacy, and the New Technology. In Lyon, D. and Zureik, E. *Computers, Surveillance, & Privacy* (eds). University of Minnesota Press. 1996.

27. Mancini, C. Thomas, K., Rogers, Y., Price, B.A., Jedrzejczyk, L., Bandara, A. K., Joinson, A. N. and Nuseibeh, B. 2009. From spaces to places: emerging contexts in mobile privacy. In Proc. of Ubicomp '09. ACM, New York, NY, USA, 1-10.

28. Mayer-Schönberger, V. Generational Development of Data Protection in Europe. In Agre, P. and Rotenberg, M. *Technology and privacy: the new landscape*, pp. 219-243. MIT Press. 1998.

29. Nissenbaum, H. Privacy as Contextual Integrity. Washington Law Review 79 (30), 2004.

30. Palen, L. and Dourish, P. 2003. Unpacking "privacy" for a networked world. In Proc. CHI '03. ACM, 129-136.

31. Sadeh, N. Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M. and Rao, J. 2009. Understanding and capturing people's privacy policies in a mobile social networking application. Personal and Ubiquitous Computing, 13, 6 (August 2009), 401-412.

32. Schegloff, E.A. 1972. Notes on a Conversational Practice: Formulating Place, in D. N. Sudnow (ed.), Studies in Social Interaction (New York: MacMillan, The Free Press, 1972), 75-119.

33. Shklovski, I., Vertesi, J., Troshynski, E., Dourish, P. 2009. The Commodification of Location: Dynamics of Power in Location-Based Systems. In Proc. of Ubicomp '09, 11-20.

34. Stutzman, F. and Kramer-Duffield, J. Friends Only: Examining a Privacy Enhancing Behavior in Facebook. In *Proceedings of CHI '10*. ACM Press, 1553-1562.

35. Tang, K. P., Lin, J., Hong, J.I., Siewiorek, D.P. and Sadeh, N. 2010. Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. In Proc. of Ubicomp '10. ACM, New York, NY, USA, 85-94.

36. Toch, E., Cranshaw, J., Drielsma, P.H., Tsai, J.Y., Kelley, P.G., Springfield, J., Cranor, L., Hong, J. I. and Sadeh, N. 2010. Empirical models of privacy in location sharing. In Proc. of Ubicomp '10. ACM, 129-138.

37. Troshynski, E., Lee, C. and Dourish, P. 2008. Accountabilities of presence: Reframing location-based systems. In Proc. of CHI '08, ACM, New York, 2008.

38. Tsai, J. Y., Kelley, P., Drielsma, P., Cranor, L. F., Hong, J.I. and Sadeh, N. 2009. Who's viewed you?: the impact of feedback in a mobile location-sharing application. In Proc. of CHI '09. ACM, 2003-2012.