Privacy Behaviors of Lifeloggers using Wearable Cameras

Roberto Hoyle[†]   **Robert Templeman**[†♭]   **Steven Armes**[†]
Denise Anthony[‡]   **David Crandall**[†]   **Apu Kapadia**[†]

[†]School of Informatics and Computing
Indiana University Bloomington
Bloomington, IN, USA

[‡]Dept. of Sociology
Dartmouth College
Hanover, NH, USA

[♭]Naval Surface Warfare Center
Crane Division
Crane, IN, USA

ABSTRACT

A number of wearable 'lifelogging' camera devices have been released recently, allowing consumers to capture images and other sensor data continuously from a first-person perspective. Unlike traditional cameras that are used deliberately and sporadically, lifelogging devices are always 'on' and automatically capturing images. Such features may challenge users' (and bystanders') expectations about privacy and control of image gathering and dissemination. While lifelogging cameras are growing in popularity, little is known about privacy perceptions of these devices or what kinds of privacy challenges they are likely to create.

To explore how people manage privacy in the context of lifelogging cameras, as well as which kinds of first-person images people consider 'sensitive,' we conducted an *in situ* user study ($N = 36$) in which participants wore a lifelogging device for a week, answered questionnaires about the collected images, and participated in an exit interview. Our findings indicate that: 1) some people may prefer to manage privacy through *in situ* physical control of image collection in order to avoid later burdensome review of all collected images; 2) a combination of factors including time, location, and the objects and people appearing in the photo determines its 'sensitivity;' and 3) people are concerned about the privacy of bystanders, despite reporting almost no opposition or concerns expressed by bystanders over the course of the study.

Author Keywords
Lifelogging; wearable cameras; privacy

ACM Classification Keywords
K.4.2. Social Issues; K.4.1. Public Policy Issues: Privacy

INTRODUCTION
Photography and its role in our everyday lives have changed dramatically in recent years. Instead of carrying a dedicated camera, many people now use a smartphone or tablet as their primary photo-taking device. Despite the difference in form factor, smartphones are similar to legacy cameras in that a

Figure 1: Lifelogging and wearable cameras, from left: Narrative Clip, Google Glass, and Autographer.

user points the device at a scene, composes the photo with a (virtual) viewfinder, and presses a 'shutter' button to capture the image. With this deliberate mode of image collection, it is possible for people to manage the stream of data themselves — they decide not only when to collect photos, but also which images to share and not share, and with whom. The recent emergence of wearable camera devices is promoting an entirely new mode of photography in which the camera discretely and continuously captures large quantities of opportunistic images without any action from the user. This new mode of use, a concept known as 'lifelogging,' is now widely available to consumers through devices like the Narrative Clip,[1] Autographer,[2] and soon Google Glass[3] (Figure 1).

These devices take photos at such a high rate — the Narrative Clip collects up to 120 images per hour while Autographer collects up to 360 per hour — that several thousand photos can be collected over the course of a single day. Some of these images may include subject matter considered private, including snapshots of computer monitors containing private emails or banking information, or photos of people in embarrassing situations. Given the volume of images collected each day, controlling access to these images manually (i.e., by sharing or deleting) is both labor-intensive and fraught with consequences from mistaken disclosures ("misclosures" [5]). Yet we do not understand how users of lifelogging devices perceive the privacy implications of the images they collect, or whether or how they will (try to) manage image-gathering and sharing. Such an understanding could lead to better mechanisms to help people manage their privacy, especially since

---

[1] **http://www.getnarrative.com**

[2] **http://www.autographer.com**

[3] **http://www.google.com/glass/start**

lifelogging products are already on the market and likely to become more popular over time.

While lifelogging in general has received considerable research attention (e.g., through the SenseCam conference series), researchers are only beginning to understand and address the privacy implications of imagery captured by these devices. Recent work has shown how opportunistic images generate new threats to users, such as allowing 3-D models of their environment to be surreptitiously created and enabling 'virtual theft' [28]. Complementary work includes defensive frameworks in which users define policies based on physical location, so that photos taken in predefined sensitive spaces can be recognized and then deleted or quarantined for review [27]. Researchers have also qualitatively studied reactions of bystanders to wearable camera devices [10, 20], as well as the sensitivity of lifelogging data and how it could be automatically altered to enable privacy-preserving processing [30]. However, we do not have a clear understanding of the privacy behaviors and attitudes of the 'lifeloggers' using these devices, or the kinds of images that people will perceive as sensitive, either for themselves or others. Shedding light on these issues will inform approaches to help lifeloggers better manage their privacy and respect the privacy of bystanders.

Research Questions
In this paper we focus on three key research questions:

R1: *How and why do people share pictures taken by lifelogging cameras, and which characteristics of the images and the environment make users more or less likely to share lifelogging photos?* We seek to understand the circumstances under which users decide to share or not to share their opportunistic photos with different people in their social networks, as well as the reasons behind these decisions. The answers to these questions may allow us to create new privacy management techniques that analyze photos and contextual information to inform or predict sharing decisions.

R2: *How do users of lifelogging cameras manage the flow of personal information collected by their devices?* In particular, we study whether and how often people actively alter the actual *collection* of images through actions such as turning off the phone, covering the camera, and manually pausing collection. Studying these behaviors will help us understand the degree to which users prefer controlling *image collection* rather than limiting *image dissemination*, and may suggest privacy management techniques that combine these approaches.

R3: *How do users report reactions from bystanders to the presence of lifelogging devices and how willing are lifeloggers to respect the privacy of bystanders?* We seek to understand the level of acceptability of lifelogging cameras by the general public and whether lifeloggers are receptive to technologies that may improve the privacy of bystanders. We also seek to know whether lifeloggers alter their sharing behaviors out of a sense of 'propriety' for the privacy of others.

Our Contributions
To address these research questions we conducted an *in situ* user study, in which participants ($N = 36$) actively used lifelogging cameras and then commented on lifelogging imagery taken during their real-life, daily routines. Participants 'lifelogged' with our customized devices over the course of a week and answered survey questions related to the images captured from their device. In addition we conducted in-person interviews with the participants at the end of the study.

Our findings indicate that 1) some people may prefer to manage privacy through *in situ*, physical control of image collection in order to avoid later burdensome review of all collected images; 2) the 'sensitivity' of an image is determined by a combination of factors including time, location, and the objects and people appearing in the photo; and 3) people are concerned about the privacy of bystanders, and reported no negative reactions from bystanders in this study.

RELATED WORK

Lifelogging applications
Researchers have examined lifelogging systems from a variety of perspectives [6, 18, 29], including how they may affect society [2], their potential for enhancing memory retention for people with memory handicaps [16], and their use in tracking infants' milestones [15]. Other recent work has studied how lifelogging could be used as a form of presentation management, allowing a lifelogger to monitor and cultivate his or her personal image [22], or for sharing memories with friends [9]. While some of these papers have considered privacy concerns, none have studied sharing and privacy behaviors of lifeloggers through an *in situ* user study as we do. Significant legal issues may also exist with lifelogging data, as discussed by Allen [2].

Sharing lifelogging data
Sharing private sensor data (e.g. geo-location) has been studied extensively [3, 8, 24, 31, 32], however image-based lifelogging is considerably different given the rich information contained in an image. Researchers have studied how people share their life events and images online, but most of this work has focused on intentional sharing of deliberately taken photos rather than lifelogging. Microsoft's SenseCam project [17,20] specifically examined lifelogging applications but to our knowledge never performed studies about how lifeloggers manage the information. O'Hara et al. [22] proposed having separate sensors record data for private versus public consumption. This approach is not suitable for images as separate imaging sensors would still record the same scene, while logic controlling public versus private sharing could be incorporated into a single sensor.

Other systems [1, 7] proposed decreasing the accuracy of the information collected by sensors in lifelogging applications in order to preserve privacy, but while noise can easily be added to some sensor data (e.g., geo-locations), it would be difficult to transform images while still preserving their utility. More recently, the DARKLY system [14] transforms images before they are available to an application, but it targets computer vision algorithms in which the goal is to extract features from

an image in a privacy-preserving way, not to produce an image that is still viewable by a human. While selectively blurring parts of the image is one possible approach, more work is needed to understand when and why such transformations should be applied. This paper aims to shed light on this issue.

Thomaz et al. [30] proposed a privacy saliency matrix to guide which images created by a lifelogger may pose privacy threats and should therefore be protected. Their model focuses on a specific setting related to eating behaviors in the context of a user study whereas we consider lifeloggers who share images for social reasons.

Privacy of bystanders

Denning et al. [10] studied how people respond to the presence of augmented reality wearable devices. They found that bystanders assumed these devices were used for recording, and "were predominantly split between having indifferent and negative reactions to the device." Although bystander reactions is not our focus, we report on reactions as perceived by our participants (who were wearing the cameras); our study found higher levels of acceptability, possibly because our lifelogging device was worn around the neck and appeared less obtrusive than augmented reality glasses. Our work also reports on propriety behaviors of users of lifelogging devices, who may be willing to suppress sharing images that infringe on the privacy of bystanders.

METHOD

Lifelogging System

We designed an *in situ* user study where participants wore lifelogging cameras for a week and answered questions about the first-person images that were collected. Instead of relying on commercial lifelogging camera systems, which often have design choices and parameters that would be difficult or impossible for us to modify, we instead designed and built our own simple and economical lifelogging platform. This system allowed us to collect the data we needed for our study while still addressing privacy and IRB considerations, described below. The system consisted of two parts: 1) a custom lifelogging application that ran on an Android smartphone worn around the participants' necks (Figure 2) and collected images and other sensor data at regular intervals; and 2) a back-end server that stored the images and later presented them to participants via a web-based interface.

Smartphone app

Our custom Android application took periodic photos from the rear-facing camera of a Samsung Galaxy Y smartphone. It also recorded readings from its accelerometer and magnometer, and logged GPS location, orientation, and ambient sound level.[4] This data collection occurred every five minutes. At the end of each five-minute collection period, the data was either immediately sent to our back-end server via WiFi and removed from the device, or, if no WiFi network was available at that point in time, the data was cached until the next collection period. Participants were asked not to use the phone for purposes other than the study. The smartphone app itself

---

[4]Not all of this data is analyzed in this study.

had a very simple interface as shown in Figure 2, offering participants only two options: (1) to pause data collection for a period of 15, 30, or 60 minutes; or (2) to delete the last 15, 30, or 60 minutes of recorded data. The app was programmed to collect data between the hours of 8AM–10PM.

Back-end server and web-based questionnaire

Our back-end server received and stored the image and sensor data from the smartphone, and hosted a web-based interface to have participants perform end-of-day surveys. When users logged into these surveys, they were shown the images that had been collected by their device since the last time they had logged onto the site, and were asked a series of questions.

1. First, they were asked to mark images that should be immediately deleted from the study and explain their reason for deletion. Users were free to mark any images they wished, but were specifically asked to delete images with nudity, in locations where photography was prohibited, or of people who had asked not to be recorded.

2. Second, we asked the user to code images that were too blurry or contained no useful information as "unusable."

3. Each user was shown a map of the key places that they had visited throughout the day, produced by a clustering algorithm applied to the participant's GPS traces, and asked to label each cluster with one of: "Home," "Someone else's home," "Class," "Lab," "Library," "Work," "Restaurant," "Bar," "Coffee shop," "Gym," or "Other."

4. The system then presented a list of any times and places that the participant had used the pause or delete functions on the smartphone app and asked them to explain why the interruptions had occurred.

5. Next, users were shown the photos that they had not deleted or marked as unusable and were asked to identify which images they would *not* be comfortable sharing with each of four categories of people: "Close friends and family," "Other friends and family," "Co-workers, classmates, and acquaintances," and "Everyone."

6. Once the user had iterated through all groups in Step 5, we asked survey questions seeking reasons for their sharing decisions. To make this manageable, we only asked about a random subset of images from each of four categories: images shared with everyone, images shared with no one, images shared with just one group, and images shared with 2-3 groups. For images that were not shared, we asked participants how embarrassed and how angry they would feel if the image were accidentally shared with that group, and how embarrassed and angry members of that group would feel if the image were shared accidentally with others.

Study Procedure

Recruitment

Undergraduate students on a large college campus (Indiana University Bloomington) were recruited through posters placed in common areas, postings on online university classifieds boards, and email solicitations. Subjects were screened

through a Limesurvey questionnaire that verified that the subject had lived in the United States for at least five years, was at least 18 years old, and was a current university student.

Enrollment
Screened participants were provided with informed consent and study information sheets. If they agreed with the documents, they were invited to come in person to our lab, where they signed the consent form and received a hard copy of the study procedures. Enrollments occurred on Mondays, and the participants were asked to perform the study for the rest of the week. On Friday participants collected payment, answered the end-of-day survey, and were interviewed and debriefed.

Each participant was given an Android phone (with our app pre-installed) attached to a bright red lanyard so that the phone could be worn around the neck (Figure 2). To inform bystanders about the device's lifelogging activity, each lanyard had a sticker with prominent text:

<div align="center">

photography in progress
IU research study
(photos taken every 5 minutes)

</div>

Participants were also instructed on how to use the app's interface to pause and undo recordings.

Ethical considerations
This study was approved by the Institutional Review Board at Indiana University (IRB Protocol #1305011388). Because of the nature of this study, we sought and received approval from not only our IRB but also from the General Counsel (legal counsel) at our university. In consultation with these offices, we developed a list of "Do's and Don'ts" that we gave to participants, in order to reduce potential risks arising from our study. These recommendations included that participants not wear the device in situations with an expectation of privacy (e.g., dorm rooms and dorm hallways) without the consent of others present, in areas where photography is prohibited (e.g., locker rooms or government installations), in a workplace without the employer's consent, and so on.

We tried to anticipate and mitigate a variety of legal and ethical risks. As one specific example, we were concerned that bystanders could become angry that images of themselves had already been captured by a participant's smartphone. The retroactive delete functionality of our smartphone app was designed to alleviate this concern, so that participants could easily diffuse tense situations by simply deleting the data. Participants were also given a set of business cards that they could hand out to bystanders who were curious or concerned about the study. These cards contained a description of the study and a URL to further information. If a participant was confronted, the participant could give them a card, who in turn could contact us to have their data (i.e., physical likeness) removed, without the involvement of the participant. (No bystanders actually contacted us during our study.)

End-of-day questionnaires
At the end of each collection day, participants were sent an e-mail reminder to fill out the online questions about their images (described earlier). Participants were asked to fill out
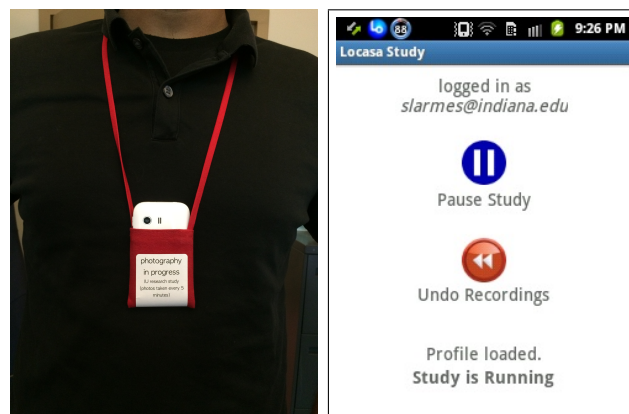

Figure 2: Our custom lifelogging platform: *(left)* smartphone and lanyard worn by participants; *(right)* screenshot of the lifelogging app running on the smartphone.

the questionnaire the first day of the study (Monday), even though it would not contain a full day's worth of data, to troubleshoot any connection issues; for our actual analysis we only used data collected from Tuesday through Thursday. If a participant missed an evening's questionnaire, the system simply had them label that day's images the next time they logged in. For survey questions about reasons for sharing or not sharing, participants could provide additional free-form responses, which we later coded in our analysis.

End-of-study questionnaires and interviews
At the end of the week, participants returned to the lab to fill out the final end-of-day questionnaire as well as an end-of-study questionnaire. This final questionnaire asked about their privacy perceptions, as well as their knowledge of lifelogging and related technologies. We also conducted an in-person interview to help us learn whether the participant ran into any issues during the study, and in particular how they interacted with their peers and how bystanders perceived them.

Compensation
Study participants were paid $50 upon completion of the study, pro-rated based on how much of the study they had completed. Participants who completed the study were entered into a drawing for an iPad Mini. (Participants were informed before the study that there would be at most 50 participants, so their chance of winning was at least 1 in 50.)

Pilot testing and study refinement
A pilot test was performed with three subjects, including two undergraduate Computer Science majors and a programmer working in our lab but not familiar with the study. The pilot study allowed us to test all aspects of the experiment, from initial screening questionnaires to the end-of-study interviews, and to find bugs and refine procedures.

FINDINGS

Participants
The study included four one-week cohorts over the course of two months in 2013 and 2014, with a total of 36 participants

Reason	Frequency
To record and remember that I had visited this place	32 (88.9%)
To share the places I visit	31 (86.1%)
To share photos of my friends	31 (86.1%)
To share what I'm doing	28 (77.8%)
To share photos of my family	27 (75.0%)
To show that I liked the place	22 (61.1%)
To appear cool and interesting	19 (52.8%)
I wanted geographically distant friends/family to feel that they were part of my day-to-day activities	18 (50.0%)
I was at a political/social/artistic event and wanted to promote it	13 (36.1%)
To demonstrate my photography skills	8 (22.2%)
To ask for information about something	5 (13.9%)

Table 1: Reasons participants gave for sharing images online.

completing the study. The participants included 21 women and 15 men, of which 21 (58.3%) were Informatics or Computer Science majors, with the remaining 14 (38.9%) having other majors and one (2.8%) undecided.

Based on their survey responses, 33 (91.7%) of our participants had previously used Facebook, 20 (56%) had used Twitter, while two (5.5%) had not used any social networking site. Among our 36 participants, 30 (83.3%) reported often sharing photos through Facebook, 12 (33%) regularly shared photos through Twitter, and 18 (50%) often used Instagram. Thirty (83.3%) had been sharing pictures on the Internet for multiple years. The most popular reasons they gave for sharing images were "to record and remember that I had visited this place," "to share the places I visit," and "to share photos of my friends." Other responses are summarized in Table 1.

Participants were relatively comfortable with others tagging them in online photos (median 5 on a 7-point Likert scale, with "1" being "very uncomfortable" and "7" being "very comfortable"). Overall, most participants were not very familiar with lifelogging services (median 3 for 'I am familiar with lifelogging or lifeblogging' on a scale with "1" being "not familiar at all" and "7" being "very familiar," with 14% responding 6 or 7).

Participants were asked a series of questions modeled on a modified Westin Scale [5], to gauge participants' preferences for and attitudes about information control, disclosure, and awareness on 7-point Likert scales (recoded so that "7" indicated high privacy). We calculated the mean responses within control, awareness, and disclosure and show the results in Figure 3. Our sample of participants strongly agree that people should be aware of, and have control over, information collection and dissemination. They had more widely varying levels of concern about sharing personal information with other parties online.

Image Collection

A total of 20,957 images were captured during the study. Images that were not reviewed by participants, or that were outside the Tuesday to Thursday study window, were removed, resulting in a total of 14,744 images included in the analyses reported here. The number of images collected per participant
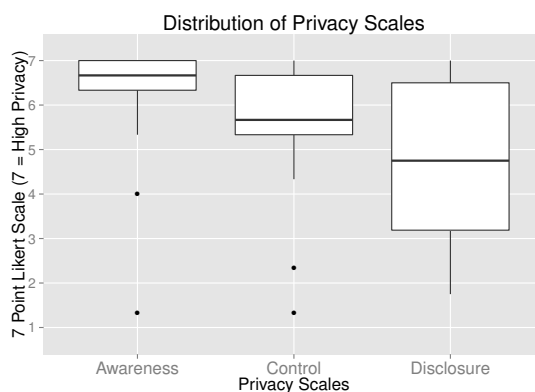

Figure 3: Privacy attitudes of our participants, by category.

ranged from 250 to 534, with a mean of 409.6 and standard deviation of 75.9, and a median of 414.5.

Participants were asked to mark 'unusable' images (e.g., dark or blurry with no identifiable objects or people) during the end-of-day questionnaires. Participants labeled between 11.8% and 100% of their images as usable (mean=20.0%, median=63.7%, standard deviation=24.0%). Of the 36 participants, 11 marked fewer than 50% of their images usable.

Managing Image Sharing
Of the 36 participants, 25 shared more than 90% of their images with all of the four sharing groups ("Close friends and family," "Other friends and family," "Coworkers, classmates, and acquaintances," and "Public"), 9 shared between 10% and 90% of their images with all four groups; and two shared fewer than 10% of their images with all four groups. In total, 92.4% of images were shared with all four groups while only 3.81% were shared with none.

These results show that most participants in our study were willing to share most images, with only a few participants actively restricting the sharing. Most participants did actively delete some images (see 'Managing Image Collection' below), but deleted images constituted a relatively small fraction of all the images collected. The low deletion rate may not be surprising since lifelogging cameras capture daily life, much of which consists of benign events. However, the number of images deleted and restricted in other ways (see below) allows us to examine patterns of both sharing and privacy-preserving behavior.

Reasons why images were shared
Table 2 summarizes why participants shared lifelogging photos, according to the end-of-day surveys. The results suggest that many people seemed to have a 'default share' approach to lifelogging, since the predominant reason given for sharing images was that 'there was no good reason to not share it,' and that most images were not considered private or sensitive. Other frequent reasons included that a particular image was 'good' or portrayed other people or the lifelogger well.

Reason	Frequency
There was no good reason to not share it	760 (82.0%)
It is a good photo	154 (16.6%)
I like how other people look / are portrayed	48 (5.2%)
I want others to see it	43 (4.6%)
I like how I look / am portrayed	27 (4.0%)
People in it would want me to share	18 (1.9%)
This picture has interesting content	13 (1.4%)

Table 2: Summary of reasons for sharing lifelogging images.

Reason	Frequency
No good reason to share it	94 (36.0%)
Objects (other than people) in the photo	80 (30.7%)
Where this photo was taken	59 (22.6%)
People within the photo	48 (18.4%)
Participant was in the photo	30 (11.5%)
It had private information	30 (11.5%)
It had uninteresting content	23 (8.8%)
It would have been embarrassing to share it	14 (5.4%)
It would have violated someone else's privacy	10 (3.8%)
It was a bad photo	4 (1.5%)
It had limited interest	4 (1.5%)

Table 3: Summary of reasons not to share lifelogging images.

Reasons why images were not shared
Table 3 summarizes the reasons people chose not to share images. Here we observe a similar 'default delete' approach, in which one-third of photos were not shared simply because there was no good reason to do so. The next most common reason for not sharing an image was due to objects (other than people) appearing in the photo. In looking at the images, we observed that these images often contained computer monitors, which may contain sensitive information on the screen. We observe that some reasons, like who was in the photo, were named both as reasons to share an image and to not share an image, which underscores the fact that the mere presence of certain people within images does not by itself signal a sensitive photo. People also reported location, as well as potential for embarrassment either to themselves or others in the image, as other reasons not to share. This latter result, combined with the fact that 'People within the photo' is another major reason not to share, suggests that lifeloggers have some degree of 'propriety' and are cognizant and sensitive to the privacy of bystanders.

Image content analysis
We manually coded some basic characteristics of the images to determine if there were obvious features that could predict whether an image was shared or not. These included whether the image was indoors or outdoors, whether there was a computer monitor in the image, and the number of people in the image. Table 4 shows the percentage of photos that were shared by participants by each type of image content.

Participants appear to be somewhat less likely to share images that were taken inside compared to outside, although this difference is not statistically significant (91.9% vs. 96.6%, $\chi^2 = 3.4$, $p = 0.07$), and the vast majority of photos were shared regardless of this location distinction. A more significant difference in sharing was seen for computer monitors: participants were much less likely to share photos with

Feature	Count	% Shared
Indoors	873	91.98%
Outdoors	120	96.67%
People present	519	92.68%
No people present	496	92.54%
Computer monitor visible	297	87.21%
No computer monitor visible	688	94.62%

Table 4: Percentage of images shared, by image content.

Shared with	Home	Work	Other
No groups	85 (1.9%)	6 (2.9%)	83 (1.8%)
Some groups	254 (5.5%)	6 (2.9%)	134 (2.9%)
All groups	4247 (93%)	193 (94%)	4437 (95%)

Table 5: Number of images shared, by location type.

a monitor compared to images without (87.2% vs. 94.6%, $\chi^2 = 16.3$, $p < 0.01$). However, even here, the majority of photos that included computer monitors were shared. Finally, we examined whether there were differences in sharing based on the presence of people in the image, but found no statistical (or substantive) difference (92.5% vs. 92.3%, $\chi^2 = 0.01$, $p = 0.93$).

Location-based analysis
We identified a total of 534 geo-location clusters by analyzing the GPS traces of each participant. As described above, we asked participants to label these clusters. Table 5 presents photo sharing rates according to location type, grouped into three categories: "home," "work," and "other." We see that participants shared a large number of photos with all groups regardless of location, suggesting that most images captured in particular locations are not necessarily considered sensitive or private. Participants shared the vast majority of photos regardless of location, though they were somewhat less likely to share images taken at home compared to other locations (92.6% versus 95.3%, $\chi^2 = 29.5$, $p < .01$).

Overall, the fewest images were captured at work (2.2%) compared to home (48.5%) and other (49.3%). This is likely because our undergraduate participants spend less time in formal work situations, but may also be because they used physical discipline at work to avoid capturing images at all.

Understanding the consequences of accidental sharing
For images that participants chose not to share, we analyzed how 'angry' or 'embarrassed' participants reported they would feel if an image were accidentally shared with a group from whom they had explicitly withheld it. To do this, we considered only the 167 images from 18 users that had reasons for not sharing *other than* that they 'had no good reason to share.' A total of 52 (31%) images were not shared because of potential embarrassment and 54 (32%) because of potential anger by the participant. For each of the 18 users who had taken one of these images, we computed the percentage of images that they chose not to share for which they would be angry/embarrassed if it were accidentally shared, and then averaged these across the 18 participants. We found that participants said they personally would experience anger for about 29% of photos on average, and would experience

embarrassment for about 22% of photos. On average, participants believed people within the photo would experience anger for about 16% of unshared photos and embarrassment for about 19% of images, suggesting that at least some participants did not share photos for reasons of propriety, that is, the concern for the privacy of bystanders.

Managing Image Collection

Most participants were willing to share most of the images that were recorded by their lifelogging devices. However, participants were able to prevent photos from being taken in the first place, because our lifelogging application had a simple user interface that could either pause data collection or retroactively delete recent blocks of collected data. We find that most participants used *in situ* pause and delete events more than post-collection deletion (Figure 4), showing a preference for limiting *collection* of images in certain situations rather than reviewing (and possibly deleting) images later.

As described above in Methods, the end-of-day web survey asked participants to provide reasons for their use of these pause and delete features. We describe our findings about these control events in the following subsections.

In-situ pause events
Participants used the pause feature to proactively interrupt recording (i.e., before they were in a context in which recording was not desired or permitted) between 0 and 19 times over the 3-day study, with an average of 6.33 or about twice per day. Six participants did not use the pause functionality at all. The leading reasons for pausing data collection were: using the bathroom (34.8%), being in a location or situation where recording is prohibited (35.4%), avoiding photographing someone nearby (5.4%), and because of nearby nudity (5.2%). The last two reasons show propriety behavior — concern for the privacy of others.

In-situ delete events
The retroactive delete function of our mobile app allowed participants to immediately delete recently-collected data, in case participants did not anticipate a sensitive situation ahead of time or forgot to pause collection but remembered later. Only seven of 36 participants used this functionality at all, and these participants used the feature between one and five times. The leading reasons for invoking the retroactive delete function were using the bathroom (38.1%), because of nearby nudity (23.8%), and being in a location or situation where recording is prohibited (14.3%).

End-of-day deletion of images
Participants were also asked to review images during their end-of-day surveys, and delete ones they wanted permanently removed from the study. Participants deleted a total of 476 images through this interface, although this total was heavily skewed by a single participant who deleted 214 images. Sixteen participants deleted no images at all, and the remaining 19 users deleted a mean of 13.8 images (median of 4).

Participants were asked to provide one or more reasons for deleting images including a freeform 'Other' field. The majority of responses (282) did not include a reason for why the
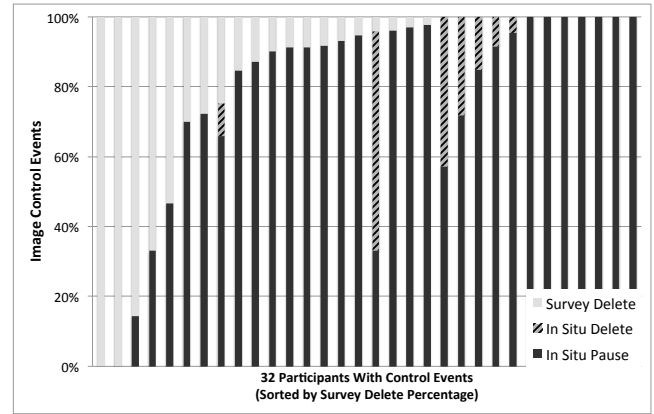

Figure 4: Frequency of control exerted by each participant, through *in situ* pause or deletes and end-of-day surveys.

Reason	Frequency
No Reason	282 (59.2%)
Unusable Picture	158 (33.2%)
Computer Monitor / Screen	10 (2.1%)
People in Photo	10 (2.1%)
Participant was in Photo	8 (1.7%)
Photo Location	3 (0.6%)
Objects in Photo	2 (0.4%)
Nudity	2 (0.4%)
Participant Error	2 (0.4%)

Table 6: Summary of why photos were deleted during end-of-day review. Participants could provide more than one reason per image. (Total number of deleted images = 476.)

image was deleted. The next largest category (158) indicated an unusable image. Participants had been asked not to delete unusable photos at this stage, but a large number ignored that request. As shown in Table 6, the reasons given for 34 of the remaining 36 photos were based on content within the images: people (10 responses), computer monitors/screens (10), participants (7), location (3), other objects (2), and nudity (2). The final two were due to participant error.

Deletions by location
Table 7 shows the type of physical privacy control used in different locations and that participants used these controls differently across locations ($\chi^2 = 27.1, df = 2, p < 0.01$). Participants were more likely to use the in situ pause when in locations other than home (60.6% in other vs. 45.0% in home), while they were more likely to use the survey delete at home than in other locations (35.0% in other vs. 54.2% in home). Participants used the in situ delete least often of the three types of physical privacy control, but were more likely to use it in locations other than home (4.4% in other vs. 0.8% in home).

Other modifications of behavior
Participants reported during the post-study interviews that they preferred to take the device off when entering the bathroom and other locations where they were not supposed to record, rather than to delete the images after the fact. A majority (23) of our participants reported that wearing the lifel-

	Home	Other
In situ pause	112 (45.0%)	248 (60.6%)
In situ delete	2 (0.8%)	18 (4.4%)
Survey Delete	135 (54.2%)	143 (35.0%)

Table 7: Control events by location.

ogging device did not cause them to modify their behavior, so for the most part they attended the same activities that they would normally attend. Of those that reported modifying their behavior during the study, most reported that they avoided some places/activities in order not to record people that they thought would not want to be recorded, or not to record behavior that they did not want recorded. On the other hand, two participants reported purposely engaging in different behavior than usual because they wanted to record it.

Privacy of Bystanders
Our end-of-study questionnaires and interviews were used to gauge reactions of bystanders to lifeloggers in our study. We found that bystanders were generally accepting of the lifelogging technology, and that lifeloggers themselves engaged in propriety behaviors to help protect bystander privacy.

Bystander reactions reported by participants
Participants reported various types of reactions from bystanders. One participant reported that their roommate was uncomfortable with the camera and avoided interaction for the duration of the study. Others reported that some of their friends would pose in front of the device for pictures. Several participants commented that they got curious stares from bystanders in public. One participant reported getting many questions in private settings from friends and acquaintances, but no questions or confrontations in public settings.

Participants reported using different mechanisms for getting consent from bystanders. In one class, the professor asked the students to vote on whether the device was to be allowed to record, and the class agreed. Two participants reported that they got consent from their sorority house members to be able to wear the device in the house. A few participants reported that they explicitly put the device away when interacting with people who did not want to be photographed.

Given the media attention to privacy issues related to devices like Google Glass and the Narrative Clip [12, 19, 21], we were surprised to find how seemingly uncontroversial lifelogging was among the participants. No participant reported encountering conflicts with bystanders or other opposition to the study during the interviews, while 26 out of our 36 participants reported positive interest from bystanders. A majority (24) of participants reported that nobody asked them to pause or turn off the device, or to otherwise not record them. Half of the requests to pause recording came in class or at work.

Propriety attitudes of lifeloggers
Participants were asked whether they agreed or not (with a 7-point scale from "strongly disagree" being "1" to "strongly agree" being "7") with the following statements relating to consent for lifelogging devices:
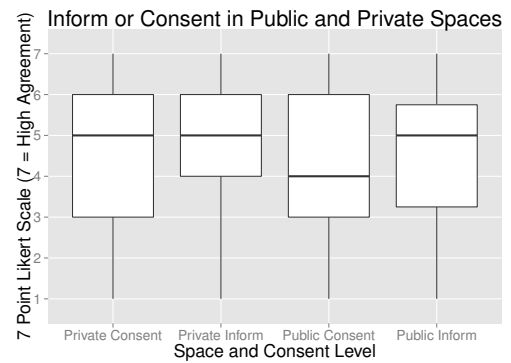

Figure 5: Participants' attitudes about collecting consent or informing bystanders of lifelogging devices.

- In *non-public* spaces, someone using a lifelogging device should have the *consent* of bystanders to include their photos in lifelogs.
- In *non-public* spaces, it is important that bystanders in the vicinity of someone using a lifelogging device be *informed* as to what the lifelogging device is doing.
- In *public* spaces, someone using a lifelogging device should have the *consent* of bystanders to include their photos in lifelogs.
- In *public* spaces, it is important that bystanders in the vicinity of someone using a lifelogging device be *informed* as to what the lifelogging device is doing.

Their responses are summarized in Figure 5. For private spaces, participants agreed on average that bystander consent is needed (61% of participants were in agreement, response of 5, 6, or 7 on a Likert scale), but many participants also did not agree (36%, response of 1, 2, or 3 on a Likert scale). Most also agreed that bystanders should be informed of the lifelogging device (67% in agreement and 22% against). In public spaces participants were mixed about whether bystander consent was needed (40% in agreement, 46% against) but generally agreed that bystanders should be informed (56% in agreement, 26% against). In general the study participants agreed that bystanders should be informed about lifelogging devices. The observed sharing and delete behaviors presented earlier in Findings support the view that lifeloggers engage in (or support) propriety in honoring bystanders' privacy.

Comfort with the Lifelogging Experience
We asked participants to rate their comfort with the following services on a Likert scale ('1' being 'Extremely Uncomfortable' and '7' being 'Extremely Comfortable'):

- An 'always on' system where pictures are automatically and constantly made available to those they have authorized.
- Services where they share pictures (taken by themselves or others) deliberately by uploading pictures they want to share and authorizing explicitly who is able to view them.

Figure 6: Participants' comfort with services where lifelogging photos would be shared automatically versus manually.

The results in Figure 6 show that there is a statistically significant (paired t-test, $t = 5.9, df = 35, p < 0.01$) difference in participants' comfort levels between the 'always-on' nature of lifelogging as compared to services where photos are uploaded deliberately and sporadically, showing that participants were less comfortable with an 'always-on' system.

We also asked participants to report their own comfort (on same 7-item Likert scale) using the logging device during the study. A majority (64%) reported comfort (comfort rating ≥ 5) using the device (median 6), including in both public (median 5, with 61% comfortable) and private (median 5, with 67% comfortable) environments. They also expressed comfort in being around someone else using a lifelogging device (median 6, with 61% comfortable), though their desire to purchase and use a device was low (median 2, with 31% expressing desire to purchase). Though participants wanted to participate in a study like this, it was not because they had a strong interest in lifelogging devices or applications, because few had heard of or were familiar with any existing devices aside from Google Glass.

DISCUSSION AND IMPLICATIONS

Various factors make a photo sensitive
A major motivation behind our study is to understand whether, when, and why people choose to share and not share images taken from their lifelogging cameras. This is an important step towards our eventual goal of creating automatic mechanisms to detect images that capture private or sensitive information that users may not want to share.

We found that no single factor seemed to determine when or why users chose not to share (or to delete) images. Many participants shared by default, and chose not to share simply when there was no good reason to share an image. A few participants had the opposite preference of deleting most images. In either case, the reasoning seemed to be similar: for the large numbers of images having no real interesting subject matter, some users may have opted to share them anyway (since there was no harm in doing so) or to not share them (because there was no use in sharing them). In the future we plan to examine these images to better understand this category of behavior. Other reasons participants gave for either deleting or not sharing indicated that *objects* (including computer monitors), *location*, and *people in the photo* (including the participant) influenced their decisions. Previous

work related to location sharing shows that privacy is typically context-dependent; in other words, it is a combination of who, what, where, and when rather than any one simple dimension [3, 13, 23].

Our work is a first step towards uncovering the various ways lifeloggers think about their privacy in the context of images. Further work is required to develop automatic algorithms that might be able to precisely identify what makes an image sensitive or private (e.g., as features in a machine-learning classifier). We draw three specific implications from our results:

- Sensitive information is often displayed on computer monitors, which are often captured by lifelogging devices. An obvious first step for automated algorithms should be to detect images with computer monitors or other types of device screens.

- Image sharing behaviors were similar across locations. Further analysis is needed on more fine-grained location labels, but our results suggest that sensitive images are deleted through *in situ* controls and review, reducing the need to automatically identify location as a dominant factor in determining image sharing preferences. Previous work proposed algorithms to detect photos taken in specific 'sensitive spaces' [27], but more research is needed to determine how automated mechanisms can incorporate location attributes. Indoor versus outdoor photo classification has been studied in the computer vision literature and may provide a starting point for automated image analysis [4] as our results indicate people share fewer indoor images (but a large fraction are shared nevertheless).

- The identities of people appearing in a photo clearly matter, but the context of the situation and the appearance of the people need to be taken in to account. Automated mechanisms could attempt to extract how many people are in the photo, whether certain people are in the photo, whether their expressions are embarrassing, and the activity occurring in the photo. Existing work in face recognition [26] and pose detection [11] could potentially be applied here to produce automated analysis algorithms.

Lifelogging devices need usable in situ controls
A major finding of our study is that participants used physical control of the device as a primary technique to manage their privacy. Participants would either pause photo collection, engage the *in situ* (on-device) delete function, or put the device away. The high frequency of unusable photos may indicate that the device was placed in a pocket rather than just paused, or that the participant otherwise covered the device, for example underneath their clothing. Only some participants elected to actively filter a large number of images by reviewing photos through the web survey. Additionally, a large fraction of participants were willing to share most photos not controlled by their *in situ* discipline with all groups in their social network (including to "everyone").

We draw the following implications:

- The prospect of reviewing images from a lifelogging device may be viewed as cumbersome and overwhelming

to many users. One way for users to alleviate the burden of manual review is to adapt their behavior to control the collection of images consciously, instead of later trying to find and remove problematic photos. The design of lifelogging cameras should thus provide adequate on-device controls such as those made available in our study.

- We suggest the following specific controls: 1) it should be easy to pause and resume the collection of images; 2) it should be easy to retroactively delete images, since people cannot always anticipate sensitive situations; and 3) it should be easy to physically remove, obscure, or cover the camera, since users may not have time to pause collection via software, and may not trust software controls in highly sensitive situations. Some form factors may make this difficult; for example if Google Glass is equipped with prescription lenses, it may not be feasible to take the glasses off for a visit to the bathroom.

- It is unclear if people will be able or willing to alter their behavior permanently to physically control their wearable cameras, especially when they might have to control multiple wearable devices (such as smartphones and fitness trackers). Automated mechanisms could relieve the burden of physical control by automating context-dependent decisions. Barring such mechanisms, however, users might prefer physical control.

Lifeloggers care about the privacy of bystanders
Our study found that one of the reasons that participants gave for deleting photos was because of the people present in the photo. Thus an interesting and important implication is that the people wearing lifelogging cameras actually care about the privacy of bystanders and actively try to delete or not share photos of them.

In other words, our results suggest that people using lifelogging cameras may be willing to specify 'propriety preferences' — situations under which they are willing to discard, modify, or not share images when they may violate the privacy of others. A similar concept has been proposed in the context of location sharing [25]. If images that might violate bystander privacy could be recognized automatically, lifeloggers could specify propriety policies to limit sharing them. We hypothesize that if the use of such propriety policies is widespread (e.g., enabled by default on lifelogging services), lifelogging in general may be more accepted.

Lifelogging cameras are acceptable to bystanders
Our study participants described mostly positive experiences with bystanders. Most people were curious about the lifelogging device and were comfortable with the device once they understood how photos were being collected (e.g., that speech was not being recorded). Participants reported that their friends and acquaintances usually exhibited interest and potential concern with the device, but strangers rarely posed questions. It is possible that some of the bystanders were less concerned about the collection because the photos would be used in a research study, but in general bystanders did not express concerns to study participants, while friends said they thought the device was interesting and 'cool.'

LIMITATIONS
This study was conducted with 36 undergraduate college students at one university over the period of one week. None of the participants had used a lifelogging device prior to the study. Our findings should thus be considered exploratory and seen as a first step towards uncovering people's behaviors and attitudes in the context of camera-based lifelogging. Further studies focusing on different populations in different stages of their lives are needed for a fully generalizable result.

Legal issues involved in studying a student population will also affect the ecological validity of the findings. Due to legal concerns, our participants were explicitly forbidden from wearing their devices in many locations in which they could have been recording on their own volition. Outside of an academic study, however, even Google has produced a list of "Do's and Don'ts" for using Glass,[5] attempting to defuse some tensions that have been felt by bystanders.

Finally, while we do report on bystander reactions, these reactions were from the perspective of our participants wearing the cameras; we did not interview the bystanders themselves.

CONCLUSION
While newly available wearable camera devices offer exciting functionality, their ability to capture large volumes of images impacts the privacy of lifeloggers as well as bystanders captured in such images. Through an *in situ* study we shed light on how people use and perceive such devices in the context of managing their privacy, and we explore some of the factors that may contribute to whether users think an image is private. Our findings motivate and provide grounding for further research into devising automated mechanisms to detect sensitive images based on these and other factors.

We show that lifeloggers are concerned about the privacy of bystanders and actively limit the dissemination of images that may impact them. While techniques akin to digital rights management (DRM) have been suggested as a way for bystanders to limit collection and dissemination of images with their likeness, we believe these approaches will be unrealistic, since lifeloggers may be unwilling to cede control of their device. Our findings instead motivate an interesting socio-technical approach based on 'propriety settings,' which if used by enough lifeloggers may reduce (though not eliminate) the privacy concerns of bystanders.

---

[5] `https://sites.google.com/site/glasscomms/glass-explorers`

REFERENCES

1. Al-Muhtadi, J., Ranganathan, A., Campbell, R., and Mickunas, M. Cerberus: a context-aware security scheme for smart spaces. In *Proceedings of the First IEEE Conference on Pervasive Computing and Communications* (2003).

2. Allen, A. Dredging up the past: Lifelogging, memory, and surveillance. *The University of Chicago Law Review* (2008), 47–74.

3. Anthony, D., Henderson, T., and Kotz, D. Privacy in location-aware computing environments. *IEEE Pervasive Computing 6*, 4 (2007).

4. Boutell, M., and Luo, J. Beyond pixels: Exploiting camera metadata for photo classification. *Pattern Recognition 38*, 6 (June 2005), 935–946.

5. Caine, K. *Exploring everyday privacy behaviors and misclosures*. PhD thesis, Georgia Institute of Technology, 2009.

6. Chen, Y., and Jones, G. Augmenting human memory using personal lifelogs. In *Proceedings of the 1st Augmented Human International Conference* (2010).

7. Cheng, W., Golubchik, L., and Kay, D. Total recall: are privacy changes inevitable? In *Proceedings of the the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences* (2004).

8. Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., and Powledge, P. Location disclosure to social relations: Why, when, & what people want to share. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2005).

9. Crete-Nishihata, M., Baecker, R., Massimi, M., Ptak, D., Campigotto, R., Kaufman, L., Brickman, A., Turner, G., Steinerman, J., and Black, S. Reconstructing the past: personal memory technologies are not just personal and not just for memory. *Human–Computer Interaction 27*, 1-2 (2012).

10. Denning, T., Dehlawi, Z., and Kohno, T. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *International Conference on Human Factors in Computing Systems* (2014).

11. Duan, K., Batra, D., and Crandall, D. A Multi-layer Composite Model for Human Pose Estimation. In *Proceedings of the British Machine Vision Conference (BMVC)* (2012).

12. Financial Times - A candid camera that stalked me in the bathroom, Jan. 2014. `http://www.ft.com/cms/s/0/d44fc9d6-8756-11e3-ba87-00144feab7de.html` (Accessed June 6, 2014).

13. Iachello, G., Smith, I., Consolvo, S., Chen, M., and Abowd, G. D. Developing privacy guidelines for social location disclosure applications and services. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)* (2005), 65–76.

14. Jana, S., Narayanan, A., and Shmatikov, V. A Scanner Darkly: Protecting user privacy from perceptual applications. In *Proceedings of the 34th IEEE Symposium on Security & Privacy* (2013).

15. Kientz, J., Arriaga, R., and Abowd, G. Baby steps: Evaluation of a system to support record-keeping for parents of young children. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009).

16. Lee, M., and Dey, A. Lifelogging memory appliance for people with episodic memory impairment. In *Proceedings of the 10th International Conference on Ubiquitous Computing (UbiComp)* (2008).

17. Mann, S., Nolan, J., and Wellman, B. Sousveillance: Inventing and using wearable computing devices for data collection in surveillance environments. *Surveillance & Society 1*, 3 (2003).

18. Miluzzo, E., Lane, N., Fodor, K., Peterson, R., Lu, H., Musolesi, M., Eisenman, S., Zheng, X., and Campbell, A. Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys)* (2008).

19. New York Times - Shine Starts to Wear Off a Little for Google Glass, May 2013. `http://bits.blogs.nytimes.com/2013/05/03/the-shine-starts-wears-off-google-glass/` (Accessed June 6, 2014).

20. Nguyen, D., Marcu, G., Hayes, G., Truong, K., Scott, J., Langheinrich, M., and Roduner, C. Encountering SenseCam: Personal recording technologies in everyday life. In *Proceedings of the 11th International Conference on Ubiquitous Computing (UbiComp)* (2009).

21. NPR All Tech Considered - Cool Or Creepy? A Clip-On Camera Can Capture Every Moment, Feb. 2014. `http://www.npr.org/blogs/alltechconsidered/2014/02/24/280733721/cool-or-creepy-a-clip-on-camera-can-capture-every-moment` (Accessed June 6, 2014).

22. O'Hara, K., Tuffield, M., and Shadbolt, N. Lifelogging: Privacy and empowerment with memories for life. *Identity in the Information Society 1*, 1 (2008), 155–172.

23. Patil, S., Gall, Y. L., Lee, A. J., , and Kapadia, A. My privacy policy: Exploring end-user specification of free-form location access rules. In *Proceedings of the Workshop on Usable Security (USEC)*, vol. 7398 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg (Feb. 2012), 86–97.

24. Patil, S., Norcie, G., Kapadia, A., and Lee, A. Reasons, rewards, regrets: Privacy considerations in location sharing as an interactive practice. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS)* (2012).

25. Reilly, D., Dearman, D., Ha, V., Smith, I., and Inkpen, K. "Need to know:" examining information need in location discourse. In *Proceedings of the 4th International Conference on Pervasive Computing (PERVASIVE)* (2006).

26. Taigman, Y., Yang, M., Ranzato, M., and Wolf, L. Deepface: Closing the gap to human-level performance in face verification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014).

27. Templeman, R., Korayem, M., Crandall, D., and Kapadia, A. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *21st Annual Network and Distributed System Security Symposium (NDSS)* (2014).

28. Templeman, R., Rahman, Z., Crandall, D., and Kapadia, A. PlaceRaider: Virtual theft in physical spaces with smartphones. In *Proceedings of The 20th Annual Network and Distributed System Security Symposium (NDSS)* (2013).

29. Teraoka, T. Organization and exploration of heterogeneous personal data collected in daily life. *Human-centric Computing and Information Sciences 2*, 1 (2012).

30. Thomaz, E., Parnami, A., Bidwell, J., Essa, I., and Abowd, G. Technological approaches for addressing privacy concerns when recognizing eating behaviors with wearable cameras. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)* (2013), 739–748.

31. Toch, E., Cranshaw, J., Hankes-Drielsma, P., Springfield, J., Kelley, P. G., Cranor, L., Hong, J., and Sadeh, N. Locaccino: A privacy-centric location sharing application. In *Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing*, UbiComp '10 Adjunct (2010).

32. Wang, Y., Leon, P., Scott, K., Chen, X., Acquisti, A., and Cranor, L. Privacy nudges for social media: An exploratory facebook study. In *Proceedings of the International Conference on World Wide Web (WWW)* (2013).