Reflection or Action?: How Feedback and Control Affect Location Sharing Decisions*

Sameer Patil[♯,†] **Roman Schlegel**[♯,‡] **Apu Kapadia**[♯] **Adam J. Lee**[§]
[♯]School of Informatics and Computing, Indiana University, Bloomington, IN 47408, USA
[§]Department of Computer Science, University of Pittsburgh, Pittsburgh, PA 15260, USA
[†]Helsinki Institute for Information Technology HIIT, Aalto University, 00076 Aalto, Finland
[‡]Corporate Research, ABB Switzerland Ltd., CH-5405 Baden-Dättwil, Switzerland
sameer.patil@hiit.fi, roman.schlegel@ch.abb.com, kapadia@indiana.edu, adamlee@cs.pitt.edu

ABSTRACT
Owing to the ever-expanding size of social and professional networks, it is becoming cumbersome for individuals to configure information disclosure settings. We used location sharing systems to unpack the nature of discrepancies between a person's disclosure settings and contextual choices. We conducted an experience sampling study ($N = 35$) to examine various factors contributing to such divergence. We found that *immediate* feedback about disclosures without any ability to control the disclosures evoked feelings of oversharing. Moreover, deviation from specified settings did not always signal privacy violation; it was just as likely that settings prevented information disclosure considered permissible *in situ*. We suggest making feedback more *actionable* or delaying it sufficiently to avoid a knee-jerk reaction. Our findings also make the case for proactive techniques for detecting potential mismatches and recommending adjustments to disclosure settings, as well as selective control when sharing location with socially distant recipients and visiting atypical locations.

INTRODUCTION
The ubiquity of portable devices equipped with WiFi and cellular network access has made it easy to track devices' locations and has provided people with many ways to share location information with others. For instance, many people use location sharing systems (LSS) to share their location and 'geotag' data. Such systems may be standalone (e.g., Foursquare) or embedded within other platforms (e.g., Facebook Places). Recent studies have shown that people find utility in sharing location information for a variety of reasons, including meeting up with friends, receiving discounts and other rewards, promoting oneself, and increasing the involvement of distant relations in everyday life [22].

*A large part of this research was conducted while the first two authors were researchers at Indiana University.

CHI 2014, April 26–May 1, 2014, Toronto, ON, Canada.

http://dx.doi.org/10.1145/2556288.2557121

Despite their utility, the use of LSS is not without risks. LSS such as Google Latitude and Apple Find My Friends allow contacts to monitor each other's locations without explicit action by the person being tracked. Such an approach raises privacy concerns because people may not anticipate possible sensitive situations in advance (e.g., a visit to the mental health clinic). Disclosure preferences to control location sharing may not be applicable when handling such unforeseen conditions. To address this limitation 'in-the-loop' systems require users to grant or deny each location request explicitly. In-the-loop systems are necessarily more disruptive to the people being queried. To strike a balance between online and in-the-loop modes, researchers have proposed ways to convey 'exposure feedback' that raises awareness and allows users to reflect on accesses to their data without requiring interaction for each and every request [23, 24, 26].

Research Questions
Our research aims to advance the usability of exposure aware LSS by better understanding *when* and *why* individuals' in situ choices diverge from their disclosure settings. To this end, we address the following research questions:

R1: To what degree do a priori location disclosure preferences differ from in situ sharing decisions and why? Disclosure settings supported by most LSS require projecting the rich context of everyday life onto relatively few dimensions (e.g., identity of requester, day, time). Even though these settings are imprecise, it is unclear whether the imprecision could be systematically characterized based on contextual factors.

R2: How do people react to immediate feedback of location exposure (i.e., a notification that one's location has been disclosed) as compared to making in-the-loop disclosure decisions? In other words, how do people feel when the system uses their pre-specified access-control rules to make a decision on their behalf and how does this compare to the burden of being asked to approve each individual location request?

R3: Are unusual (e.g., infrequently visited) locations more likely to cause disagreement between a person's disclosure preferences and in-context decisions? If so, to what degree?

R4: What other contextual factors (e.g., relationship to requester, strength of the relationship, current activity, etc.)

contribute to disagreements between pre-specified disclosure settings and in situ decisions?

Studying these aspects can provide insight into effective ways of conveying exposure feedback via a deeper understanding of situations in which users have heightened location sharing concerns compared to what they previously anticipated. Properly addressing the above research questions requires in situ feedback from individuals engaged in real-life scenarios. Therefore, we chose the Experience Sampling Method (ESM) to conduct a study in which participants ($N = 35$) used our 'Locasa' LSS to respond in situ to hypothetical requests made by members of their social circles.

Contributions

Our findings highlight two key points. First, although prior work has shown that reflecting on exposure feedback increases LSS utility and user comfort [23, 24, 26], our results indicate that *immediate* feedback without the opportunity to control actions actually increases user *discomfort* with location sharing and leads to feelings of oversharing. Second, we note that deviation from an individual's specified settings may not always signal a privacy violation: in our study, around half of the time, deviations were the result of *withholding* information that the individual would have been comfortable disclosing, indicating a potential loss of LSS benefit.

RELATED WORK

Several areas of related work are relevant to our study.

ESM Studies of Location Sharing Behavior

Several researchers have applied ESM [15] to study location sharing privacy. Consolvo et al. [6] report the results of an ESM study assessing the behavior of 16 participants in a simulated LSS. This study found that the three most important factors influencing the willingness to share location were the identity of the requester, the specificity of location description, and the reason for requesting access. Khalil and Connelly [12] conducted an ESM study with 20 participants to examine how sharing information on location and other activities was impacted by social relationships and contextual conditions. Anthony et al. [1] conducted an ESM study with 25 undergraduates to explore the relationship between *place* (broadly defined as physical location plus social context) and the willingness to share location, finding that sharing decisions were often based on more than mere physical location. These works studied contextual factors that impact location sharing decisions. In contrast, we investigated how such decisions are affected by exposure feedback and control.

Privacy Policies for Location Sharing

As noted above, common factors that influence location sharing include physical location, relationships, and context [1, 6, 12]. Although difficult to pin down precisely, these factors are typically captured in LSS that support identity-, time-, and/or location-based rules [23, 26]. In addition, others have noted that privacy preferences could be predicated on the activity in which an individual is engaged [11], the individual's perception regarding the utility to the requester [4], the entropy of the current location [25], and the frequency with which location has been requested by others [24]. Efforts have been made to develop policy languages or systems capable of capturing these types of conditions to enable better end user control of location privacy [7, 9, 16, 20, 23, 26]. Further, Benisch et al. [3] collected and analyzed location disclosure preferences of 27 individuals over 3 weeks to determine trade-offs involved in accurately matching users' intentions vs. reducing the burden of specification. We studied how feedback and control affect sharing decisions and expect these findings to apply to several such policy constructions.

We employed a policy model in which sharing rules are based on recipient groups, time, and day. Although relatively simple, this model is largely consistent with research prototypes (e.g., [23, 26]) as well as widely deployed LSS, such as Foursquare, Google Latitude, and Facebook Places.

Privacy Expectations in Context-Aware Systems

Social media users often unintentionally reveal private information. For example, Mao et al. [18] showed how people on Twitter revealed information related to vacation plans (which could lead to one's home being burglarized [5]), health conditions, and even drunk driving. Researchers have also found that people may regret sharing information on social networks [22, 27]. Among other reasons, people have reported concerns related to possible job loss due to the location revealed; indeed, there have been several media reports of individuals losing jobs as a result of Facebook postings [8]. While previous works highlighted privacy leaks and surveyed users about regrets associated with social media and LSS, our work sought to uncover people's privacy expectations during actual use of such systems.

Feedback and Exposure Control

Researchers have studied mechanisms for providing feedback to help raise user awareness of potential inconsistencies with a priori preferences. Many techniques have been investigated, including an 'audit log' of exposure detailing access histories in IM [10] and LSS [26], interactive feedback to guide policy authoring [21], and ambient feedback on smartphone home screens to convey the frequency of location accesses [24].

Although explicit feedback can help users control their level of exposure, it is well known that frequent or inopportune interruptions can be counterproductive and distract users from their primary goals [2]. Tsai et al. [26] studied the *usefulness* of exposure feedback and found that being able to view a log of prior location accesses was helpful to users. However, the study did not explore the extent of disagreement between the users' original disclosure settings and those after auditing the log. Sadeh et al. [23] focused on continual policy refinement via machine learning as well as feedback based on user inspection of location disclosure history. They showed that such refinement can indeed improve privacy. However, they noted that feedback utility "plateaued" after a point, beyond which in-context decisions were not suitably handled by their system.

Our goal was to gain insight into effective and non-intrusive feedback techniques that are likely to help people manage their exposure and avoid regrettable disclosures. Toward this

end, we aimed to characterize the situations in which people's sharing choices are likely to diverge from their previously specified sharing policies. We also explored people's reactions to exposure feedback related to automated sharing decisions (based on a priori preferences).

METHOD

To conduct the study, we built an Android app called "LocasaESM" that leverages our location gathering infrastructure called Locasa. At several random times throughout the day, the app presented participants with a questionnaire based on hypothetical location inquiries made at that moment by members of their social network.

Locasa

LocasaESM collected locations at 15-minute intervals using only the cellular and wireless networks (i.e., no GPS). Our internal pre-study testing showed that these settings provided a reasonable tradeoff between battery life and location resolution.

Questionnaires in the app could be triggered at relevant moments based on the scheme we specified. Participants were informed of a questionnaire via a sound as well as a notification icon in the phone's status bar. Tapping the notification brought up the questionnaire. To ensure timely sampling of experience, questionnaires disappeared if not answered within 15 minutes.

Recruitment and Participants

We recruited participants from our institutional neighborhoods to be able to provide in-person assistance or troubleshooting, if needed. We advertised in communities at and near two large public US universities: one in a college town (Bloomington, IN) and one in a metropolis (Pittsburgh, PA). Inclusion of two different kinds of localities provided greater sample diversity. In Bloomington, we advertised in the 'Job' category of an online campus bulletin board accessible to the entire university community. In Pittsburgh, we recruited via a university-maintained participant pool that included individuals from across the city. To increase participant diversity beyond university populations, we advertised in the 'Et Cetera jobs' category of each town's Craigslist. To avoid priming, the study was advertised without revealing the privacy focus. Participation was spread over February to April 2013.

Study Procedure

Screening

The advertisement directed potential participants to a brief screening questionnaire. Given the influence of culture on privacy, we sought to minimize cultural diversity of the sample by restricting participation to those over the age of 18 who indicated having lived in the US for at least 5 years.[1] Since LocasaESM ran only on Android phones, we further limited participation to Android users with cellular data access.

[1]Prior research indicates that sufficient cultural assimilation can be assumed after 5 years [13].

Enrollment

Those who met the screening criteria were presented with an online consent form detailing the study procedures. Upon consent, participants created an account in the Locasa system and proceeded to the following steps:

Initial access-control setup. We sought to improve upon previous relevant studies, such as the one by Tsai et al. [26], that set only global rules for all recipients. Therefore, participants were asked to name four location recipients in each of the following categories: Family, Friends, Colleagues/Peers, and Acquaintances. Participants were further asked to indicate their relationship with each of these recipients and rate the closeness of the relationship with a five-point Likert item (from 'Not close' to 'Very close'). By grouping people into categories for which privacy preferences were likely to be similar, we sought to reduce the time and tedium of specifying separate privacy preferences for each recipient, while providing finer-grained control than possible with a single monolithic set of access permissions. We selected the four recipient groups to cover a mix of social and professional contexts as well as a wide spectrum of social and/or professional distance. They also include the common types of individuals with whom people share location information. By including four recipients per category we hoped to increase the potential range of attitudes within each category and reduce specification of access permissions targeted at a single recipient.

Next, participants created access-control rules to specify when recipients from each group were allowed access to their location. We avoided conflicts among rules by requiring that rules be created only for allowing (and not for denying) location. Participants could create any number of rules based on day(s) and times(s). During times not covered by any of the specified rules, location requests were denied. We did not provide the ability to create location based rules in order to avoid complexity and burden and guard against participants dropping out during signup due to the time and effort of rule specification. Typical commercial and academic LSS also do not provide location based rules. Moreover, the sets of locations visited by our participants were not known to us in advance. To ensure a stable set of rules for the duration of the study, we disallowed rule changes once initial setup was completed.

App installation. Next, participants followed instructions to install LocasaESM on their phones and were randomly assigned to one of the two study conditions described below.

ESM questionnaires

During the study, we simulated location requests from the 16 recipients listed by each participant. The access-control rules created during setup were then used to determine whether location would have been disclosed or withheld for each of these requests. The requests were distributed randomly throughout each day of the 15-day study period. We chose a period of 15 days to capture weekday and weekend routines as well as to increase the likelihood of collecting non-routine locations. A longer period would have been desirable but would likely have posed difficulties in recruiting partic-

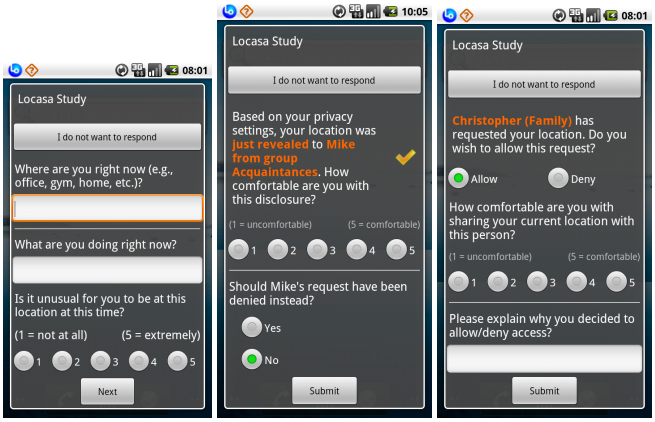

Figure 1. ESM questionnaire screenshots: The image on the left shows Screen 1 for both conditions. The images on the right show Screen 2 for Feedback and Decision conditions respectively.

ipants willing to tolerate multiple interruptions per day for more than half of a month.

Study conditions (Feedback and Decision). Each day LocasaESM presented participants with simulated location requests along with corresponding brief two-screen questionnaires. To maximize data collection while minimizing interruptions, we chose to present five requests at random between 8AM and 9PM with an interval of 2–3.5 hours between requests. To ensure late evening coverage (9PM–12AM), an additional questionnaire was presented if the phone was detected to be in active use during this period.

The first questionnaire screen (Figure 1, left) asked the participant to describe where he/she was, what he/she was doing, and whether this location/activity was unusual (on a five-point Likert item from 'Not at all' to 'Extremely'). The second screen differed across the two study conditions:

Feedback condition (Figure 1, middle) aimed to shed light on attitudes toward system decisions based on pre-specified settings *after* these decisions were made. In this condition, the questionnaire informed the participant that his/her location was sought by one of the 16 recipients he/she specified. The participant was also informed whether the location was disclosed or withheld based on the access rules he/she created at the beginning of the study. The participant could indicate his/her comfort with the disclosure with a five-point Likert item (from 'Uncomfortable' to 'Comfortable'). The participant could also indicate whether location should have been disclosed instead of being withheld or vice versa and provide an open ended explanation.

Decision condition (Figure 1, right) allowed us to examine how closely preferences expressed at the beginning of the study matched decisions made 'in-the-loop' by asking the participant to *mediate the location sharing decision*. In this condition the questionnaire alerted the participant that one of the 16 recipients was requesting location access. The participant then chose whether access should be permitted or declined and provided an open-ended explanation. If access was allowed, the participant could further indicate comfort with the disclosure with a five-point Likert item (from 'Uncomfortable' to 'Comfortable').

Post-study questionnaire and interview

The day after the conclusion of the study, participants were sent a link to a post-study questionnaire that asked about experience with social media and location sharing, collected responses for scales on consumer and interpersonal privacy, and gathered demographics. We also conducted 10–15 minute semi-structured interviews with participants willing to be interviewed. The goal of the interviews was to verify the effectiveness and validity of study operation, and gather feedback for improving the Locasa system and procedures for future studies. Two interviews were conducted in person and the rest by phone.

Compensation

We used an engagement based payment to ensure continued participation. Participants were paid $3.30 for installing LocasaESM, $0.15 for answering a questionnaire, $0.50 for answering *all* questionnaires on a given day, $3.00 for answering the post-study questionnaire, and $5.00 for providing post-study interview comments. Participants could choose to be compensated with an Amazon gift certificate or cash (paid in person or via Paypal).

Pilot testing and study refinement

All study aspects, including instructions and setup, went through several iterations of evaluation and internal testing. We also conducted a pilot with several external evaluators.

FINDINGS

Fifty one individuals qualified for the study and completed enrollment by installing LocasaESM. Of these, we excluded 15 because they did not participate for the entire study duration and/or did not complete the post-study questionnaire. We further excluded one participant who indicated in post-study comments that he did not always answer the questionnaires truthfully due to tedium and interruption. We report findings based on responses of the remaining 35 participants, 19 in the Feedback condition (Bloomington: 12 and Pittsburgh: 7) and 16 in the Decision condition (Bloomington: 10 and Pittsburgh: 6). Although a majority of the sample comprised undergraduate and graduate students (aged 19–24), they spanned diverse fields, such as Fine Arts, Community Health, Finance, and Biology. Moreover, many of these student participants held jobs outside of their studies. We also managed to recruit non-students; at least 10 participants were older than 24, with 3 older than 45.[2] Overall, participants answered 2,034 questionnaires (Feedback: 1,095 and Decision: 939), with a mean of 58 and median of 61 questionnaires per participant. Twenty six participants were interviewed at the end of the study.

Figure 2 shows a heat map combining the a priori disclosure rules of all participants for the four recipient categories. The redness of each time slot is proportional to the number of participants who allowed location access during that slot.

[2]We have ages for only 23 out of the 35 participants.

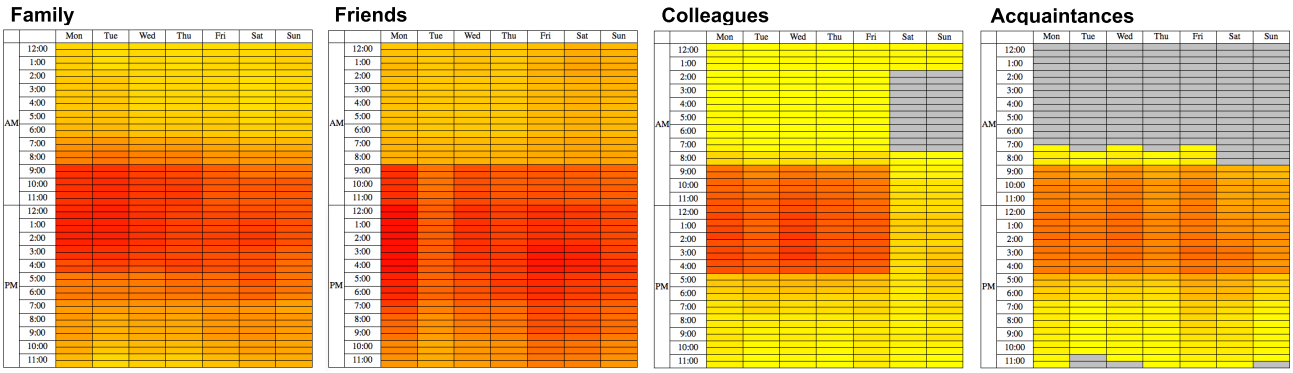

Figure 2. Heat maps showing aggregate temporal preferences of the participants for sharing location with different recipient categories.

The heat map revealed differences in access rules across categories and time periods as well as similarities in desires across participants.

Undersharing vs. Oversharing

Participants in the Feedback condition were informed of the disclosure decision made by Locasa's application of their pre-specified disclosure rules and asked for their *in situ* agreement with the outcome. Participants in the Decision condition were asked to make a decision in situ *without being informed of the decision that would have been made based on their pre-specified rules*. We compared participant responses — agreement with pre-specified preference in the Feedback condition and the actual decision made in the Decision condition — to the decision reached by the pre-specified preference. Each instance where a participant's in situ desires differed from pre-specified rules was marked as a 'mismatch.' Mismatches were classified into two categories: 'undersharing' occurred when participants wished to reveal location even though it was withheld by pre-specified rules while 'oversharing' occurred when participants expressed the desire to withhold location although their rules indicated otherwise.

Since the *occurrence* of a mismatch and the *type* of mismatch are binary variables, we employed binomial logistic regressions using each of these as the outcome variable. The former indicated whether the in situ choice was a mismatch with the preference rules, while the latter examined whether the mismatch resulted in undersharing or oversharing. Study condition (Feedback or Decision), location context (such as unusualness and recipient category), and participant characteristics were employed as predictor variables. We tested for interaction among the variables and included the participant ID as a random effect to account for repeated measures. To obtain the most parsimonious models, we examined initial results and removed predictor variables and higher-order interactions that did not exhibit statistically significant effects. Table 1 shows the resulting regression models for the two outcome variables, viz., mismatch occurrence and mismatch type. Binomial logistic regression examines each level of the predictor variable *separately*, against one of the levels treated as a baseline. Therefore, we also report results of an ANOVA of the *overall effect* of each predictor variable, along with the corresponding F-statistic. Note that the models in Table 1

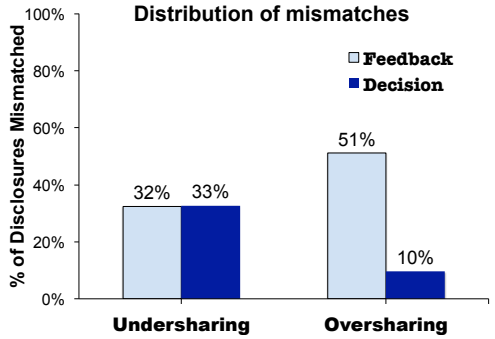

Figure 3. Distribution of mismatch types across study conditions as a fraction of *total* disclosures (mismatched as well as matched).

represent effects that simultaneously take all predictor variables into account. In the following subsections, we unpack the impact of each of the variables separately.

Match between a priori preferences and in situ decisions

Across both conditions, nearly 65% of the 2,034 responses resulted in a mismatch. Half of these (N = 662) were a result of undersharing, with the remaining half (N = 650) resulting from oversharing. The types of mismatches, however, were not similarly distributed across conditions (see Figure 3). It can be seen that the distribution of undersharing and oversharing was the opposite for the two study conditions. Further, undersharing stayed at roughly the same levels regardless of study condition while oversharing was much lower in the Decision condition.

These results suggest that *immediate feedback of location exposure evoked greater feelings of oversharing*. However, when participants were in control of making the decision and were not reminded of specified rules, oversharing was much smaller (10% compared to 51% in the Feedback condition). In contrast, the levels of undersharing were almost identical in both study conditions. Binomial logistic regression confirmed that differences between the study conditions were statistically significant for occurrence as well as type of mismatch (see Table 1). The extent of undersharing suggests that *initial settings were overly protective*, leading to location being withheld even when participants were comfortable disclosing it.

Factor	Mismatch Occurrence (Mismatch or No mismatch)					Mismatch Type (Undersharing or Oversharing)				
	Odds ratio	95% CI	p	F	Pr(F)> p	Odds ratio	95% CI	p	F	Pr(F)> p
Study condition (baseline: Condition = Feedback)				F(1, 2018)=472.449	< 0.001***				F(1, 1296)=228.086	< 0.001***
Condition = Decision	0.111	(0.046,0.270)	< 0.001***			0.007	(0.002, 0.026)	< 0.001***		
Location recipient (baseline: Recipient = Family)				F(3, 2018)=1.545	0.201				F(3, 1296)=50.116	< 0.001***
Recipient = Friends	1.589	(0.947, 2.666)	0.078			1.092	(0.622, 1.918)	0.760		
Recipient = Colleagues	1.126	(0.687, 1.844)	0.638			0.077	(0.045, 0.130)	< 0.001***		
Recipient = Acquaintances	0.778	(0.488, 1.243)	0.294			0.063	(0.037, 0.107)	< 0.001***		
Location unusualness (baseline: Unusualness = 1)				F(4, 2018)=9.54	< 0.001***				F(4, 1296)=1.509	< 0.001***
Unusualness = 2	0.917		0.776			0.751	(0.420, 1.344)	0.335		
Unusualness = 3	1.206	(0.645, 2.256)	0.558			0.738	(0.410, 1.328)	0.311		
Unusualness = 4	0.551	(0.309, 0.983)	0.044*			0.568	(0.301, 1.071)	0.081		
Unusualness = 5	0.625	(0.314, 1.246)	0.182			1.188	(0.582, 2.427)	0.636		
Study condition : Location recipient				F(3, 2018)=3.989	0.008**				F(3, 1296)=35.212	< 0.001***
Condition = Decision : Recipient = Friends	0.860	(0.446, 1.659)	0.008**			1.344	(0.483, 3.739)	0.571		
Condition = Decision : Recipient = Colleagues			0.643			38.676	(14.724,101.596)	< 0.001***		
Condition = Decision : Recipient = Acquaintances	1.064	(0.573, 1.977)	0.845			54.266	(19.989,147.322)	< 0.001***		
Study condition: Location unusualness				F(4, 2018)=8.348	< 0.001***				F(4, 1296)=13.398	< 0.001***
Condition = Decision : Unusualness = 2	1.227	(0.581, 2.593)	0.592			5.910	(2.178 ,16.038)	< 0.001***		
Condition = Decision : Unusualness = 3	0.810	(0.358, 1.835)	0.613			3.820	(1.200, 12.158)	0.023*		
Condition = Decision : Unusualness = 4	1.652	(0.711, 3.838)	0.243			5.303	(1.536, 18.304)	0.008**		
Condition = Decision : Unusualness = 5	2.079	(0.839, 5.154)	0.114			0.969	(0.247, 3.807)	0.965		

Statistical significance: *** p < 0.001, ** p < 0.01, * p < 0.05

Table 1. Results of binomial logistic regressions for 'mismatch occurrence' and 'mismatch type.'

Impact of unusualness and type of locations

Participant ratings for unusualness of locations covered the whole spectrum from 'Not unusual'(1) through 'Extremely unusual' (5). The greatest proportion (56%) of locations at which participants received the questionnaires were rated at the lowest level of unusualness. We hypothesized that mismatches would increase with level of unusualness since access rules typically cover routine and expected scenarios. However, a deeper look at the distribution of mismatches across unusualness ratings (see Figure 4) showed that the distributions were not linear; in both conditions extremely unusual locations appeared to deviate significantly from the trend in the case of undersharing as well as oversharing. We expected that preferences specified a priori would lead to high levels of oversharing in highly unusual, i.e., non-routine, places. This appeared to be valid in the Feedback condition. Interestingly, the opposite was the case in the Decision condition, where almost all mismatches at highly unusual locations resulted from participants choosing to disclose their location in opposition to their rules. This suggests that *disclosures of highly unusual locations are perhaps more readily acceptable when the sharing decision is made explicitly (Decision condition) as opposed to when feedback of a system-made decision is conveyed with no recourse to influencing the outcome (Feedback condition).*

Further examining Figure 4 revealed that the distributions of undersharing and oversharing were the opposite in the two conditions; undersharing was higher and oversharing was lower in the Feedback condition, while the opposite was the case in the Decision condition. However, across both conditions, the *differences* between undersharing and oversharing levels were the largest at extremely unusual as well as not unusual locations. In the Feedback condition, the gap between the levels of undersharing and oversharing decreased with increase in unusualness, with the exception of extremely unusual locations. This trend suggests that *the feelings of oversharing evoked by exposure feedback may be tempered for deviations from expected practices, as long as the deviations are not extreme*.

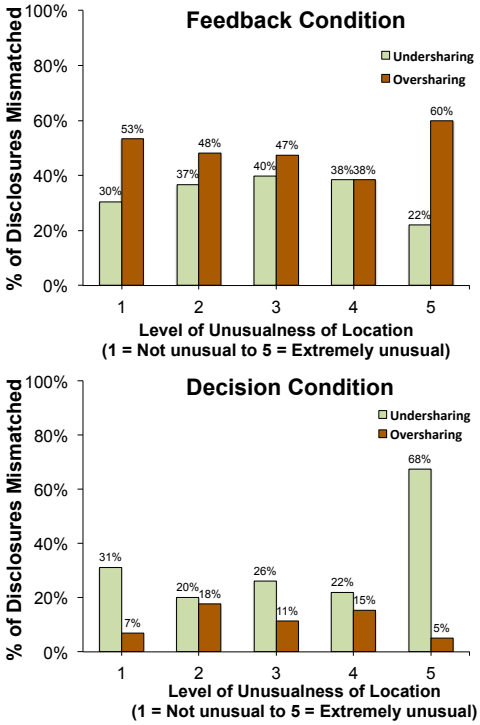

Figure 4. Mismatches across different levels of unusualness as a fraction of *total* disclosures (mismatched as well as matched).

It should be noted that that *unusualness was not tied (merely) to specific physical locations*. All participants who were interviewed confirmed that they assigned unusualness ratings based on the likelihood of being at that location at that time; the same location could differ in unusualness depending on context. The freeform location descriptions entered by the participants were independently coded by four coders into 23 common location types, such as home, office, restaurant, bar, car, and so on. After resolving coding discrepancies, the five most common locations were: home, office, someone else's home, classroom, and campus. The unusualness ratings as-

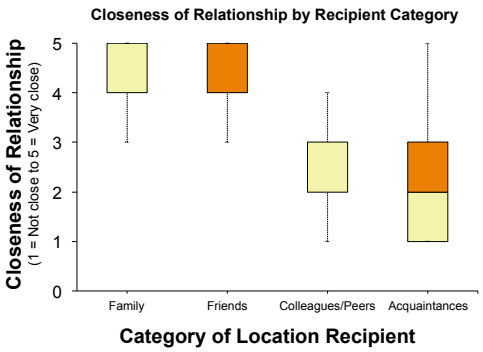

Figure 5. Boxplots showing closeness of relationship with different categories of location recipients.

signed to these five locations covered the whole range from 1 through 5.

Impact of relationship with requester

When considering mismatches combined across conditions, we found that the occurrence of mismatch (regardless of type) did not vary by relationship with the location recipient (i.e., family, friends, colleague/peers, or acquaintances). However, when examining the type of mismatch, the differences between the categories of location recipients were statistically significant. We found that colleagues/peers and acquaintances differed from family and friends. Figure 5 confirms that, compared with family and friends, colleagues/peers and acquaintances were rated to be socially distant. We expected that access rules would more likely result in oversharing with socially distant recipient categories (colleagues and acquaintances). Contrary to those expectations, we found that *participants were much more willing to share locations with colleagues and acquaintances than their rules allowed*. This was reflected in the greater percentage of undersharing than oversharing with colleagues and acquaintances (see Figure 6). Figure 2 suggests that this could be due to more cautious initial permissions set for colleagues and acquaintances compared to family and friends.

Recipient categories further exhibited statistically significant interaction effects with the study condition. Firstly, we found that the extent of undersharing with colleagues and acquaintances relative to the extent of oversharing was larger in the Decision condition (see Figure 6). Moreover, for family and friends the types of mismatches were distributed in opposite ways in the two conditions. This suggests that *the effect of feedback in evoking feelings of oversharing may be more salient in the case of family and friends*.

Impact of other contextual factors

Freeform text reasons entered by participants suggest that other contextual factors also played a role in location disclosure decisions. For instance, being tired, being occupied in an activity, or being at a place that the recipient disapproved were some of the reasons offered for denial of location access in opposition to specified rules. Post-study interview comments revealed that recency and/or frequency of access were also taken into account. For instance, one participant, who

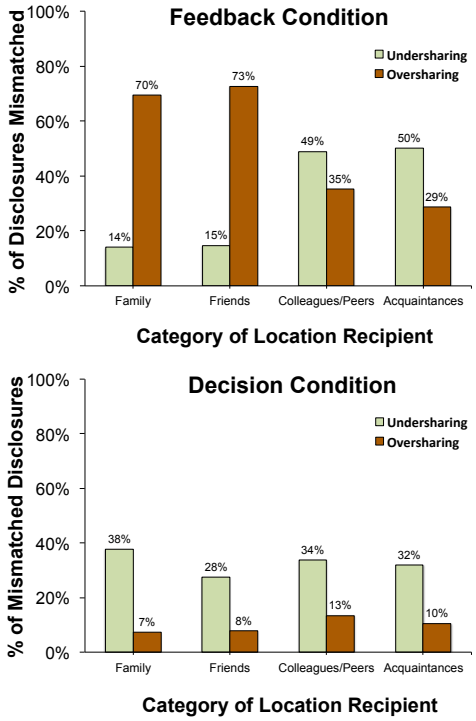

Figure 6. Mismatches across different categories of location recipients as a fraction of *total* disclosures (mismatched as well as matched).

was concerned due to too many repeated requests, said that *"It would be cool to set the number of times people from a group can request your location in a given day."* Note that disclosure settings in current systems do not typically incorporate these aspects.

We did not find any relationship between mismatches and participants' concerns regarding consumer privacy measured using the Internet User's Information Privacy Concern (IUIPC) scale [17]. One reason was lack of sufficient variance; most participants reported high levels of consumer privacy concerns (mean = 5.9, median = 6.0, on a seven-point scale with 7 being the highest level of concern). It may also be the case that the nature of privacy concerns in consumer and interpersonal domains is relatively distinct. We did note that differences in the levels of interpersonal privacy concern toward various social groups exhibited some impact on mismatches. However, these levels were unevenly distributed across scales and study conditions, so we could not include them in regressions of the whole data set.

DISCUSSION AND IMPLICATIONS

We carefully considered external validity since our data was obtained from hypothetical ESM. Such data can indeed yield useful generalizable insight as indicated by the utilization of hypothetical requests and scenarios by a large number of published location sharing studies (e.g., [4,14,23,24,28]), including those that employed ESM [1,6]. Moreover, the influence of simulation was evenly distributed due to the random allocation of participants to the two study conditions, leaving comparative analyses unaffected. To further ensure response

validity and accuracy, despite the 'fake' nature of location requests, we tactfully probed participants during post-study interviews. All interviewees indicated that they strived to treat the location requests as real. This is reflected in remarks such as *"I know you are doing the study for a purpose. It's not just a game. So I tried to give you pretty accurate information."* and *"I tried to keep it as realistic as possible about how I would feel if they were actually contacting me."* Some participants further commented that the reflection made them more aware of online information disclosure in general and location disclosure in particular. For instance, one remarked *"You think somebody is a coworker only and then once you have these alerts and it's asking whether this is appropriate, it makes you think if it really isn't appropriate or it is fine and you had set it not to be fine."* while another commented *"I know it's not actually shared but I still felt like my privacy was invaded when it was hypothetically shared with someone when I didn't really want them to know where I was."* We can therefore be fairly confident that study responses reflect real-life preferences.

Moreover, simulation provided several advantages compared to studying user behavior in deployed systems. Firstly, the investigation could focus on the core topic (i.e., location sharing) without influence of system-specific functionalities and affordances. For example, in the Foursquare LSS, gamification and economic incentives affect location sharing practices. Secondly, not being restricted to specific systems reduced participant selection bias; people could participate in our study regardless of which LSS they used. Thirdly, specific systems, especially those that do not enjoy widespread adoption, would not have adequately covered a person's social and professional networks. Yet, broad coverage was crucial when studying practices intertwined with these networks. Fourthly, 'in the wild' usage does not allow easy control over the diversity, volume, and pace of data collection. For instance, only four location requests per participant occurred during the two weeks of field deployment of the Locyoution system [26], whereas we gathered data regarding nearly 60 requests per participant over the same duration.

Location-sharing preferences

The temporal and social patterns observed in participants' access rules (see Figure 2) were in line with expectations, suggesting that participants did put thought into their rules. It is still conceivable that the time and effort needed to create access rules for four different recipient categories led to underspecification of true preferences. Figure 2 further highlights commonalities in disclosure preferences across individuals that could be useful for constructing default profiles. For instance, location disclosures during business hours on weekdays seem acceptable across recipient categories.

Mechanisms to reduce mismatches

We provided participants with finer access-control mechanisms than a single set of global privacy preferences typical of most systems. Participants could share location differently with different social groups and further control sharing according to day and time. Yet, nearly 65% of the responses suggested a mismatch between preferences and practice. Since decisions were made during the 15 days immediately after specifying preferences, it is unlikely that these mismatches could be attributed to change of circumstances after specification (e.g., breakup, graduation, job change, etc.). No major life-altering events were mentioned by interviewees. Such a high discrepancy therefore calls not only for more effective interfaces for preference specification, but, more importantly, for techniques that automatically detect potential mismatches and elicit in-the-loop reflection and action when appropriate. It should be noted that our participants could not refine their initial preferences during the study and were not informed that they would be unable to edit initially created rules. However, post-study interviews indicated that a majority of participants did not feel the need to edit their preferences during the study. Some interviewees did indicate that they would have liked to refine initial rules upon gaining more experience using the system, suggesting that LSS could consider providing preference adjustment mechanisms similar to Mazurek et al.'s [19] 'reactive access control' for shared files.

Immediate vs. delayed exposure feedback

Participants in the Feedback condition exhibited many more disagreements with a priori sharing rules than those in the Disclosure condition. Although we are unsure why in-context feedback evoked so many disagreements, the findings are a cautionary note about emotional reaction to immediate yet *after the fact* feedback *without decision-making autonomy and control*. In situations where disagreements are likely, it might be more effective to provide feedback *prior to* an impending disclosure, with ability to 'veto' the decision made by access rules. In other cases, feedback could be delivered in the aggregate (instead of each individual disclosure) and after a cooling-off period (instead of instantaneously)—e.g., [23, 24, 26].

Undersharing of location

We were surprised by the relatively high proportion of cases in which participant settings undershared location. We suspect that people initially picked conservative disclosure settings to avoid oversharing and, as a result, ended up undersharing. Undersharing could also be due to the 'opt-in' nature of Locasa's access controls; location disclosures were permitted only upon explicit user specification, which required time, effort, and thought. It would be interesting to conduct a similar study where disclosure is 'opt out' and examine whether the relative proportions of mismatch types are reversed.

An examination of freeform text responses indicates that hard-to-quantify contexts influenced participants' desires to share beyond a priori preferences. This was reflected in comments such as *"I am in the town next to her if she wants to know where I am to come hang out," "becoming a close friend, I don't mind her knowing,"* and *"I just talked to my parents; they are dropping something off."* This observation has two design implications. To date, the majority of attention to incorrect sharing preferences has centered around situations in which individuals *overexposed* information in social media (e.g., [5, 8, 18, 22, 27]). In order to enable people to derive better utility from social media, it might be worth exploring whether techniques developed to control information

overexposure could be adapted to also deal with *underexposure*. Another alternative is to allow requesters to ask the target individual to make an in-the-loop decision for a contextually motivated location request (e.g., face-to-face collaboration) that might be denied by default.

A compromise could be hybrid sharing systems that selectively prompt for in-the-loop decisions in exceptional circumstances. For instance, our findings revealed that participants were more likely to indicate that their settings undershared location with colleagues and acquaintances. Further, the likelihood of accepting disclosures contrary to settings may be higher when requesters are in physical proximity (*"I am in the town next to her..."*), or users are at specific locations (*"he could come over if he knew I was home"*). It may be possible to increase the precision of rule based sharing decisions by leveraging in-the-loop decisions in these types of contextually meaningful situations.

Unusualness of location

As discussed earlier, 'highly unusual' locations may predict potential mismatches with peoples' pre-specified preferences. Future systems could thus attempt to detect major deviations from routine. Detecting highly unusual locations could allow an LSS to suppress automated sharing decisions that may alarm users (participants were most concerned about oversharing when immediate feedback was provided in highly unusual locations) in favor of facilitating in-the-loop decisions (when given control over the disclosure choice, participants were much more willing to share location in highly unusual situations). Very unusual locations may also signal *additional willingness* to share location (e.g., when visiting a new art exhibit). Future systems must thus help users manage both facets of exposure, viz., enhancing desired exposure as well curtailing undesired exposure.

LIMITATIONS

Although we strived for breadth and diversity, a large fraction of our participants were students. Therefore, our sample cannot be considered representative of the broader US population. We note, however, that students do provide at least two benefits. First, university students tend to have less structured routines and thus higher variability in daily activities. This variance provides more opportunities for observing the impact of unexpected situations on sharing choices. Second, younger age groups are heavy social media users, thus interesting in their own right. Participants reported occasional cases in which simulation resulted in queries from recipients who were either co-present with the participant or already knew the participant's location. Such scenarios would not arise in a field trial within a deployed LSS.

CONCLUSION

We reported on an experience sampling study aimed at uncovering mismatches between pre-specified access preferences and in situ decisions for disclosing location. We found a notably high amount of mismatches, which highlights the need for more effective interfaces for specifying access permissions as well as better mechanisms for detecting contextual conditions that lead to disagreement with a priori sharing preferences. We identified commonalities in access preferences and various contextual factors that could aid in these endeavors. Interestingly, we found that providing feedback immediately after system-enacted disclosures may create a heightened sense of disagreement with the decision. It may therefore be advisable to delay disclosure feedback, allowing for a reasonable 'cooling off' period. Our findings also hint that oftentimes typical privacy preferences might not be relevant for highly unusual places; giving people in-the-loop control over decisions in such locations could *increase* location sharing. Future work can explore techniques for detecting such locations based on data from smartphone sensors.

ACKNOWLEDGMENTS

We thank the study participants and those who helped test and pilot the study. Thanks to Zhipeng Tian for help with parts of the study implementation and Roberto Hoyle for helping code location descriptions. Bart Knijnenburg provided invaluable input on statistical analyses. The paper benefited from insightful comments of anonymous reviewers. John McCurley and Sara E. Justinger provided editorial input on the final draft. This research is supported by NSF grants CNS-1016603 & CNS-1017229, and US DHS grant no. 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P). The contents of this paper do not necessarily reflect the views of the sponsors.

REFERENCES

1. Anthony, D., Henderson, T., and Kotz, D. Privacy in location-aware computing environments. *IEEE Pervasive Computing 6*, 4 (2007), 64–72.
2. Bailey, B. P., and Iqbal, S. T. Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. *ACM Transactions on Compututer-Human Interaction 14*, 4 (2008), 1–28.
3. Benisch, M., Kelley, P., Sadeh, N., and Cranor, L. Capturing location-privacy preferences: Quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing 15*, 7 (2011), 679–694.
4. Biehl, J. T., Rieffel, E. G., and Lee, A. J. When privacy and utility are in harmony: Towards better design of presence technologies. *Personal and Ubiquitous Computing 17*, 3 (2013), 503–518.
5. Bilton, N. Burglars said to have picked houses based on Facebook updates, Sept. 2012. `http://bits.blogs.nytimes.com/2010/09/12/burglars-picked-houses-based-on-facebook-updates/` (Accessed May 29, 2013).
6. Consolvo, S., Smith, I. E., 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*, CHI '05, ACM (New York, NY, USA, 2005), 81–90.
7. Cuellar, J. R., Morris Jr., J. B., Mulligan, D. K., Peterson, J., and Polk, J. M. Geopriv requirements, Feb.

2004. RFC 3693. http://www.ietf.org/rfc/rfc3693.txt.

8. Facebook post gets worker fired, Mar. 2009. http://sports.espn.go.com/nfl/news/story?id=3965039 (Accessed May 29, 2013).

9. Hengartner, U., and Steenkiste, P. Implementing access control to people location information. In *Proceedings of the 9th ACM Symposium on Access Control Models and Technologies*, SACMAT '04, ACM (New York, NY, USA, 2004), 11–20.

10. Hsieh, G., Tang, K. P., Low, W. Y., and Hong, J. I. Field deployment of IMBuddy : A study of privacy control and feedback mechanisms for contextual IM. In *UbiComp 2007: Ubiquitous Computing*, vol. 4717 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg (2007), 91–108.

11. Iachello, G., Smith, I. E., Consolvo, S., Abowd, G. D., Hughes, J., Howard, J., Potter, F., Scott, J., Sohn, T., Hightower, J., and LaMarca, A. Control, deception, and communication: Evaluating the deployment of a location-enhanced messaging service. In *UbiComp 2005: Ubiquitous Computing*, vol. 3660 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg (2005), 213–231.

12. Khalil, A., and Connelly, K. Context-aware telephony: Privacy preferences and sharing patterns. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work*, CSCW '06, ACM (New York, NY, USA, 2006), 469–478.

13. Khan, R. M., and Khan, M. A. Academic sojourners, culture shock and intercultural adaptation: A trend analysis. *Studies About Languages 10* (2007), 38–46.

14. Knijnenburg, B. P., Kobsa, A., and Jin, H. Preference-based location sharing: Are more privacy options really better? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, ACM, ACM (New York, NY, USA, 2013), 2667–2676.

15. Larson, R., and Csikszentmihalyi, M. The experience sampling method. *New Directions for Methodology of Social and Behavioral Science 15* (1983), 41–56.

16. Le Gall, Y., Lee, A. J., and Kapadia, A. PlexC: A policy language for exposure control. In *Proceedings of the 17th ACM Symposium on Access Control Models and Technologies*, SACMAT '12, ACM (New York, NY, USA, 2012), 219–228.

17. Malhotra, N. K., Kim, S. S., and Agarwal, J. Internet Users' Information Privacy Concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research 15* (December 2004), 336–355.

18. Mao, H., Shuai, X., and Kapadia, A. Loose tweets: An analysis of privacy leaks on Twitter. In *Proceedings of The 2011 ACM Workshop on Privacy in the Electronic Society*, WPES '11, ACM (New York, NY, USA, Oct. 2011), 1–11.

19. Mazurek, M. L., Klemperer, P. F., Shay, R., Takabi, H., Bauer, L., and Cranor, L. F. Exploring reactive access control. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 2085–2094.

20. Myles, G., Friday, A., and Davies, N. Preserving privacy in environments with location-based applications. *IEEE Pervasive Computing 2*, 1 (Jan.-Mar. 2003), 56–64.

21. Patil, S., and Lai, J. Who gets to know what when: Configuring privacy permissions in an awareness application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '05, ACM (New York, NY, USA, 2005), 101–110.

22. Patil, S., Norcie, G., Kapadia, A., and Lee, A. J. Reasons, rewards, regrets: Privacy considerations in location sharing as an interactive practice. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, ACM (New York, NY, USA, 2012), 5:1–5:15.

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

24. Schlegel, R., Kapadia, A., and Lee, A. J. Eyeing your exposure: Quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, SOUPS '11, ACM (New York, NY, USA, 2011), 14:1–14:14.

25. Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L. F., Hong, J. I., and Sadeh, N. M. Empirical models of privacy in location sharing. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, Ubicomp '10, ACM (New York, NY, USA, 2010), 129–138.

26. Tsai, J. Y., Kelley, P. G., Drielsma, P. H., Cranor, L. F., Hong, J. I., and Sadeh, N. M. Who's viewed you?: The impact of feedback in a mobile location-sharing application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, ACM (New York, NY, USA, 2009), 2003–2012.

27. Wang, Y., Komanduri, S., Leon, P., Norcie, G., Acquisti, A., and Cranor, L. "I regretted the minute I pressed share": A qualitative study of regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, SOUPS '11, ACM (New York, NY, USA, 2011), 10:1–10:16.

28. Wiese, J., Kelley, P. G., Cranor, L. F., Dabbish, L., Hong, J. I., and Zimmerman, J. Are you close with me? Are you nearby?: Investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, Ubicomp '11, ACM (New York, NY, USA, 2011), 197–206.