Interrupt Now or Inform Later?: Comparing Immediate and Delayed Privacy Feedback *

Sameer Patil[†] **Roberto Hoyle**[♯] **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
[†]Yahoo Labs, 701 1st Avenue, Sunnyvale, CA 94089, USA
[‡]Corporate Research, ABB Switzerland Ltd., CH-5405 Baden-Dättwil, Switzerland
sameerpatil@yahoo-inc.com, rjhoyle@indiana.edu, roman.schlegel@ch.abb.com, kapadia@indiana.edu, adamlee@cs.pitt.edu

ABSTRACT
Feedback about privacy-affecting system operations is important for informed end-user privacy management. While feedback is most relevant if provided immediately, such delivery interrupts the user and risks disrupting ongoing tasks. The timing, volume, and nature of feedback is therefore critical for avoiding inopportune interruption. We varied the *timing* and *actionability* of feedback regarding accesses to a user's physical location. We found that the sense of privacy violation was heightened when feedback was immediate, but *not* actionable. While immediate and actionable feedback may sometimes be necessary, our findings suggest that moderately delayed feedback is often acceptable. A moderate delay may serve as a compromise to minimize interruption and avoid overly alarming reaction to immediate feedback. However, immediate and actionable feedback could still be beneficial when privacy sensitivity is high or ambiguous.

Author Keywords
Privacy; feedback; location sharing; experience sampling method; ESM.

INTRODUCTION AND RELATED WORK
With the growing adoption of technologies for interpersonal interaction, privacy has emerged as an important consideration. Yet, users typically find it difficult to understand and enact privacy preferences. Researchers have long argued for system feedback that helps users understand privacy-affecting system operations [1]. These operations include not just actions of the system but also interactive acts of the people with whom one interacts via the system [6, 7]. In particular, feedback increases system transparency and fosters the development of accurate mental models of its operations, thus helping users make more informed privacy decisions [6, 8].

Despite its usefulness, feedback also suffers from several disadvantages. In particular, feedback consumes cognitive resources and distracts users from ongoing tasks. Such disruption occurs even if feedback is purely informational, requiring no user action [3]. The resulting distraction grows with increasing feedback, eventually overwhelming the user. The demands on user attention are even greater when feedback is *actionable*, requiring the user to act on the information before the system can proceed. Inopportune interruptions resulting from actionable feedback can annoy users and may also lead them to pay less attention to the information or, worse, dismiss it altogether [2].

These challenges underscore that privacy feedback is most effective when delivered in the right manner at the right time at a manageable rate. Toward this end, researchers have tried to gauge user 'interruptibility' to avoid disrupting an ongoing task [5]. However, in the case of privacy feedback, knowing merely whether a user is interruptible may not suffice; for privacy-critical situations it might be necessary to interrupt at all times while for benign cases the most appropriate choice might be to avoid disruption even if the user is otherwise interruptible.

We carried out a user study in which we delivered feedback on privacy affecting system actions in two modes: *immediate* and *delayed*. Within each mode, we manipulated whether the feedback delivered information in a passive way or required active user action. Our results reveal important differences in the utility of privacy feedback in the different modes, suggesting that different modes may be useful for different privacy affecting situations. Our results motivate hybrid delivery of privacy feedback with varying levels of actionability and immediacy to deal with different kinds of privacy concerns.

METHOD
We built a prototype location gathering infrastructure called Locasa to conduct a study in which participants received feedback about hypothetical location inquiries by members

*A large part of this research was conducted while Sameer Patil and Roman Schlegel were researchers at Indiana University.

CHI 2015, April 18–23, 2015, Seoul, Republic of Korea.
Copyright 2015 ACM 978-1-4503-3145-6/15/04…$15.00
http://dx.doi.org/10.1145/2702123.2702165

of their social circle. We have previously provided detailed information and rationale regarding study design and procedures [4]. In this section, we summarize aspects of the study relevant for this Note.

Recruitment. We recruited participants from two communities: a college town (Bloomington, IN) and a metropolis (Pittsburgh, PA). Participation was spread over February to April 2013. To avoid priming, the study was advertised without revealing its privacy focus. The advertisement directed potential participants to a brief online 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. Since Locasa is an Android app, we further limited participation to Android users with cellular data access. We obtained informed consent from those who met these criteria.

Study setup. During setup, participants named four location recipients in the following categories: Family, Friends, Colleagues/Peers, and Acquaintances. Next, participants created any number of access-control rules based on day(s) and times(s) to specify when each group was allowed access to their location. During times not covered by the rules, location requests were denied. To ensure a stable set of rules for the study duration, we disallowed rule changes after setup. During the study, we simulated location requests from these 16 recipients. These hypothetical requests were distributed randomly throughout each day of the 15-day study period. For each request, the participant's access-control rules determined whether location would have been disclosed.

Study conditions. Participants were randomly assigned to one of two between-subject study conditions: 'actionable' vs. 'non-actionable' feedback.[1] Actionable feedback asked participants to decide whether to disclose their location to the requesting party. In contrast, non-actionable feedback *merely informed* participants whether location was disclosed to the requesting party based on the access-control rules specified during setup. Within each condition, participants received feedback in two ways: 'immediate' and 'delayed.' Immediate feedback was provided 5–6 random times throughout the day via Locasa's Android app (which also collected participants' locations at 15-minute intervals). On the other hand, delayed feedback was available by logging in to Locasa's Web interface, which showed location requests occurring since the previous login. For each study day, 8–10 randomly distributed location requests were generated for delayed feedback. In the non-actionable condition, delayed feedback informed participants whether their location was disclosed to the requesting party, just like immediate feedback. Since contextual actions are not possible with a delay, delayed feedback in the actionable condition asked participants about the actions they *would have* taken at the time of the location request.

Questionnaires. In all study conditions, a brief questionnaire accompanied each piece of feedback: participants were asked to describe where they were, what they were doing, and whether their location/activity was unusual (with a five-point Likert item ranging from 'Not at all' to 'Extremely'). In the non-actionable conditions, participants could indicate whether their location should have been disclosed instead of being withheld, or vice versa, and provide open-ended explanations. In the actionable conditions, participants could enter open-ended explanations for their disclosure decisions. Figure 1 shows the questionnaires for the two conditions with immediate feedback delivered to the phone. Delayed feedback conditions presented the same questionnaires via the Web, along with a map showing where the corresponding accesses occurred.

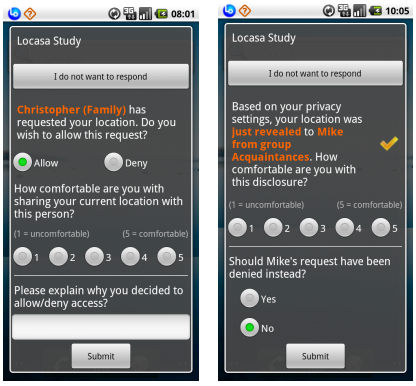

Figure 1. Immediate Feedback: Actionable (left) and Non-actionable (right).

Post-study questionnaire. At the conclusion of the 15-day study, participants completed an online 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 further conducted 10–15 minute semi-structured interviews with 26 participants who were willing to be interviewed.

Compensation. We used engagement based payments to ensure continued participation. Participants were paid $3.30 for initial setup, $0.15 for responding to a phone questionnaire, $0.50 for responding to questionnaires via the Web at the end of each day, $0.50 for answering *all* questionnaires on a given day, $3.00 for answering the post-study questionnaire, and $5.00 for the post-study interview. Participants could choose between an Amazon gift certificate and cash.

Our previous report [4] of this data was restricted to the two immediate feedback conditions delivered via the phone. This Note extends the analysis by additionally including the *delayed* versions of the respective feedback conditions.

FINDINGS

Thirty-five individuals qualified and completed the study: 19 in the non-actionable conditions and 16 in the actionable conditions. Although a majority comprised undergraduate and graduate students (aged 19–24), they spanned diverse fields of study. Many student participants were also employed part- or full-time. We did manage to recruit non-students; at least 10 participants were older than 24, with 3 older than 45.[2]

[1] 'Actionable' and 'non-actionable' feedback conditions in this Note correspond respectively with the 'decision' and 'feedback' conditions reported in the previous study [4].

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

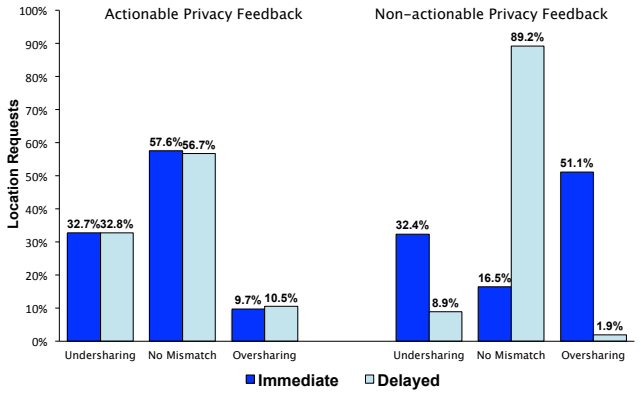

Figure 2. Mismatches between pre-specified privacy preference and contextual decision for actionable and non-actionable feedback conditions.

Across the four conditions, participants answered 5,753 questionnaires (actionable immediate: 939, actionable delayed: 1,608, non-actionable immediate: 1,095, and non-actionable delayed: 2,111). Participants in the non-actionable conditions were informed of the disclosure decision based on their access-control rules and asked for their agreement with the outcome. In contrast, participants in the actionable conditions were asked to make a disclosure decision *without being informed of the decision that would have been made by their access-control rules*.

We compared participant responses — agreement with pre-specified preferences in the non-actionable conditions and the actual decisions made in the actionable conditions — to the decisions reached by the access-control rules. Each instance where in-situ responses differed from pre-specified preferences 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 mismatch and the *type* of mismatch (i.e., undersharing or oversharing) are binary, we employed binomial logistic regressions using each of these as the outcome variable. Study conditions, location context (such as recipient category and unusualness), and participant characteristics were the predictor variables. We tested for interaction among variables and included 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.

We previously reported that for *immediate* feedback the non-actionable case evoked greater feelings of oversharing (see [4] for corresponding regression results). In this Note, we could dig deeper into this discrepancy due to the inclusion of matching study conditions with delayed feedback. Since participants in the immediate actionable condition made decisions without being informed of their initially specified privacy preference, we treated this condition as indicative of 'true' contextual desires to be employed as the comparison baseline. We then compared the baseline with the other conditions to examine the impact of the delay. As Figure 2 shows:

1. When feedback was actionable, the distribution of mismatch types stayed the same regardless of feedback timing (immediate or delayed).
2. In the case of non-actionable feedback, the feelings of oversharing almost completely disappeared when feedback was encountered after a delay.

Regression results confirmed that the distribution of mismatch type was significantly affected by feedback actionability as well as timing. There was a statistically significant interaction effect between feedback actionability and timing ($z = 8.63, p < 0.001$) indicating that mismatch distribution for non-actionable and actionable feedback was significantly different in the immediate and delayed conditions. Further, we found a statistically significant difference between immediate and delayed non-actionable feedback ($z = 11.04, p < 0.001$) but no difference for actionable feedback ($z = -0.375, p = 0.71$).

DISCUSSION AND IMPLICATIONS

As discussed in the previous section, delay affected actionable and non-actionable feedback differently. Baseline actionable feedback was unaffected by the delay. In comparison, non-actionable feedback evoked different and opposing reactions depending on the timing of delivery. When delivered immediately, it evoked a much greater sense of privacy violation. In contrast, a delay in delivery suppressed such feelings, leading to experiencing lower privacy violations than in the baseline condition. The overly alarming nature of immediate non-actionable feedback suggests that the immediate mode be reserved for the truly important and privacy-sensitive cases.

An alternative solution, of course, is to make *all* privacy feedback immediate (and actionable). The disruption resulting from such an approach makes it impractical and non-scalable. Moreover, our results show that only about 10% of requests in the baseline condition were a result of oversharing, i.e., indicative of a potential privacy violation. In other cases, immediate actionable feedback would have served no *privacy* purpose, yet would have interrupted the user. This is not to say, however, that feedback should never be immediately actionable. Such feedback is indeed necessary for avoiding oversharing. Our findings suggest several ways to keep the volume of immediate actionable feedback manageable:

- When the situation is not privacy-critical, provide feedback with a delay.
- Trigger immediate feedback at random times, particularly during initial usage when little is known about the user's contextual privacy choices.
- Utilize past decisions, along with relevant context, to detect likely cases of oversharing in order to target feedback to privacy-relevant cases.
- Examine past decisions to determine meaningful adjustments to user-specified privacy preferences that would reduce mismatches without compromising privacy.

Notably, there is a crucial framing difference between delayed feedback in the non-actionable and actionable conditions. In the case of actionable delayed feedback, participants were deciding what they would have done *at the time the location request occurred*. For non-actionable delayed feedback, however, participants answered based on their feelings about the decision *at the time of answering the question*. In other words, the reaction to non-actionable delayed feedback was influenced by hindsight and took into account whatever passed between the time of the request and the feedback. Passage of time does indeed allow various uncertainties to be resolved, thus changing prior assessments of privacy implications. In our study, passage of time typically resulted in lowering feelings of potential privacy violation.

We also point out that our findings hold methodological implications for the Experience Sampling Method (ESM). On the one hand, the lack of impact of delay in the actionable case suggests that ESM studies could conveniently gather additional data of equivalent quality via *delayed ESM*, keeping interruptions of ESM at manageable levels. On the other hand, the effect of delay on non-actionable feedback indicates that such delayed operation may be permissible when the topics involve *actions* rather than *feelings*; it seems that individuals are reasonably accurate in describing what they would have done in the recent past but find it difficult to recreate the intensity of their past feelings.

LIMITATIONS

Operational details of the study necessitated that the immediate and delayed conditions differ in terms of the feedback *interface* in addition to feedback timing. Feedback in the immediate conditions was presented via the phone while that in the delayed conditions was Web based. Therefore, it is conceivable that the results could be wholly or partially attributable to differences in the interface rather than timing. However, we believe that the contextual nature of privacy makes a strong case for attributing the findings to temporal factors.

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 and other cultures. On occasion, 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 of a deployed location sharing system.

CONCLUSION

Feedback is an important and powerful mechanism for helping users understand system operation. Such transparency is particularly useful and essential for system actions that potentially impact user privacy. The contextual nature of privacy often demands immediate attention and action. Yet, this need for immediacy must be balanced with other considerations, such as disruption and/or cognitive overload, which can potentially diminish user attentiveness and undermine the utility of privacy feedback. Using the case of location sharing, our study shows that varying the timing and actionability of feedback leads to different experiences of privacy violations. Our findings suggest that immediate feedback is best reserved for the most privacy-sensitive or ambiguous situations and potentially for calibration and learning that enables a better match between privacy preferences and in-situ system decisions.

ACKNOWLEDGMENTS

We thank the study participants. We acknowledge Bart Knijnenburg for data analysis consultation and anonymous reviewers for helpful comments. This research was supported by NSF grants CNS-1016603, CNS-1252697, CNS-1017229, and CNS-1253204, and US DHS grant no. 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P). The contents of this Note do not necessarily reflect the views of the sponsors.

REFERENCES

1. Bellotti, V., and Sellen, A. Design for privacy in ubiquitous computing environments. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work*, ECSCW (1993), 77–92.

2. Egelman, S., Cranor, L. F., and Hong, J. You've been warned: An empirical study of the effectiveness of Web browser phishing warnings. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI (2008), 1065–1074.

3. Hudson, S. E., and Smith, I. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, CSCW (1996), 248–257.

4. Patil, S., Schlegel, R., Kapadia, A., and Lee, A. J. Reflection or action?: How feedback and control affect location sharing decisions. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI (2014), 101–110.

5. Pejovic, V., and Musolesi, M. InterruptMe: Designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp (2014), 897–908.

6. Schlegel, R., Kapadia, A., and Lee, A. J. Eyeing your exposure: Quantifying and controlling information sharing for improved privacy. In *Proceedings of the Symposium on Usable Privacy and Security*, SOUPS (2011), 14:1–14:14.

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

8. Wang, Y., Leon, P. G., Scott, K., Chen, X., Acquisti, A., and Cranor, L. F. Privacy nudges for social media: An exploratory Facebook study. In *Proceedings of the 22nd International Conference on World Wide Web Companion*, WWW Companion (2013), 763–770.