Modeling End-User Affective Discomfort With Mobile App Permissions Across Physical Contexts

Yuxi Wu*†, Jacob Logas*, Devansh Ponda*,
Julia Haines‡, Jiaming Li‡, Jeffrey Nichols§, W. Keith Edwards*, Sauvik Das¶
*Georgia Institute of Technology, †Northeastern University, ‡Google, §Apple, ¶Carnegie Mellon University

***Abstract*—Users make hundreds of transactional permission decisions for smartphone applications, but these decisions persist beyond the context in which they were made. We hypothesized that user concern over permissions varies by context, e.g., that users might be more concerned about location permissions at home than work. To test our hypothesis, we ran a 44-participant, 4-week experience sampling study, asking users about their concern over specific application-permission pairs, plus their physical environment and context. We found distinguishable differences in participants' concern about permissions across locations and activities, suggesting that users might benefit from more dynamic and contextually-aware approaches to permission decision-making. However, attempts to assist users in configuring these more complex permissions should be made with the aim to reduce concern and affective discomfort—not to normalize and perpetuate this discomfort by replicating prior decisions alone.**

I. Introduction

Mobile and IoT platform permission models allow users to regulate how applications can access a variety of sensor data (including microphones, cameras, location sensing), stored data (such as files, photos, and contacts), and other resources that applications may require to provide necessary functionality. Typically, on Android, users face an ask-on-first-use (AOFU) policy, where they are prompted with permission requests when an app wants to access the data / resource gated by a permission. A user's decision under the AOFU model persists for all future uses of that app: once a permission is set, it typically persists as being always on, always off, or on only when the app is in visible use (i.e., not when running in the background) [1]. Moreover, once a user grants a permission under AOFU, they are unlikely to return to and re-evaluate the setting, even if they no longer necessarily wish to grant the permission [2], [3].

However, privacy risk and user perceptions of privacy risk vary across physical contexts [4]; do permissions decisions in one context necessarily apply to other contexts? For example, a user might expect a location permission request when on public transit, but not when they're at home. Indeed, recent work has explored how user expectations of permissions [5], as well as whether or not they would accept different permissions [6], can change across different physical locations.

In parallel, other prior work has highlighted that affective discomfort from privacy-invasive practices has become normalized in mobile app use [7]: i.e., users have come to expect and accept that mobile apps will be privacy-invasive. One way this affective discomfort is normalized is through a mismatch between extant permission models and user preferences. Given that users' perception of privacy risk, permission preferences, and permission expectations vary across physical context, we hypothesized that users' affective discomfort—i.e., feelings of concern, unease, or creepiness—towards a permissioned use of data (PUD)[1] should similarly vary across physical context (namely, user location and physical activity). Accordingly, in this work, we ask the following research question:

> **How does a user's affective discomfort with permissioned uses of data vary across physical contexts?**

To address this question, we conducted a 4-week long experience sampling study with 44 participants. Twice a day, at random times in the typical waking hours, we asked participants to report on their level of concern and attitude toward an installed application using a particular permissioned resource. We also asked participants to self-report their present context, i.e., their physical location and what they were currently doing. We then built a mixed effects linear model to describe user concern in relation to these contextual variables. As hypothesized, our findings show that users' self-reported concern over PUDs varied across physical contexts, especially when those uses conflicted with expectations. We also found differences in concern when users were in public versus private locations.

Our findings echo those of prior work (e.g., [3]) suggesting that context-aware privacy permissions are necessary to bring users' expectations of privacy in alignment with reality. But, whereas prior work has primarily explored configuring context-aware permissions on users' prior decisions, our work shows that doing so can risk further automating and normalizing users' affective discomfort. Instead, we argue that a better objective function for future permission interfaces —

[1] We distinguish between a "permission" and a "permissioned use of data" because a permission is an access control setting that applies to all physical contexts, while a permissioned use of data is the use of data enabled by that permission at a specific time, in a specific context.

Symposium on Usable Security and Privacy (USEC) 2025
24 February 2025, San Diego, CA, USA
ISBN 979-8-9919276-5-9
https://dx.doi.org/10.14722/usec.2025.23023
www.ndss-symposium.org, https://www.usablesecurity.net/USEC/

automated or not — is to measure and reduce affective discomfort with permissioned uses of data. Doing so is more likely to align with users' permission preferences than automating permission configurations based on a user's prior decisions alone, as prior work has shown that individual decisions can be made under duress or because a user feels like they have little other option (e.g., [8]).

In short, our contributions include:

- A quantitative analysis of affective responses to permissioned uses of data given physical context. We find that user concerns over PUDs vary across physical contexts.
- A brief discussion of the merits and limitations of using user affect to set permission decisions, as well as future design considerations for creating automated systems to allow and deny permissioned data use given a predicted level of concern for that use and its physical context.

II. Related Work

Permissions have been studied in-depth since the advent of modern smartphones. Past work has found that users tend to ignore permissions [9], [10]; in response, researchers have attempted—and successfully evaluated [11]—ways to address this uncertainty by making permission information more legible and accessible for users. More recently, Bonnè et al. [12] found that for common permissions—storage, phone, microphone, location, contacts, and camera—users understood why a permission was requested as they were given the context to why it is needed. A follow-up to this work [13] found that, post-Android-6.0, while users tended to reject permissions when the permission request was unexpected, they were less likely to do so when presented with an explanation for the request.

Beyond user awareness, one reason app permissions seem to be such a point of vulnerability is the mismatch between user expectations of privacy when using certain apps and in certain contexts, versus the permissions and access the apps actually have. As such, researchers have tried to pinpoint cases where user preferences and expectations of privacy clash with actual app usage. Lin et al. [14] defined a model for permissions privacy as expectation alignment, finding that when users were most surprised by an access to a sensitive resource, they also had trouble explaining why this access was necessary, and that clarifying the reason for the usage might ease user concerns. Other work has also examined whether certain application types [15] or data practices [16] can be associated with greater rates of mismatches with user expectation.

However, even a user well-educated and certain about their privacy and permission preferences can be vulnerable, through different social or psychological pressures. Users might correlate positive community ratings of apps on app marketplaces with safety, and can be misled into granting permissions [17]. Or, users might be more willing to grant app permissions when under financial duress while using mobile loan apps [8]. Outside these contexts of social and psychological pressure, users' permissions preferences could be very different.

Several attempts have been made to relieve users of overly-granular decision-making by *predicting* users' permission preferences [18], [19], [2], [1], [3]. Despite these attempts, recent research has shown that affective discomfort has been normalized in app use: users expect and accept that all apps will be "creepy." [7]. Since these attempts have primarily measured user *acceptance* of delegated decisions as the key unit of success—i.e., asking users if they would overturn the automated decisions—approaches that simply aim to learn and emulate "what a user would do" will learn to emulate behavior with this normalization of affective discomfort baked-in.

Recent work [13] has found that even users with high levels of privacy awareness can have low rates of permission denial, indicating acceptance of delegated permissions might not be a perfect match for preferences. In other words, just because a user accepts a permission setting does not mean it aligns with their privacy preferences. Indeed, researchers [7] have argued that affectively discomforting app experiences have slowly become the de facto norm for users, which allows privacy-invasive data practices to become further entrenched in the sociotechnical landscape.

We argue that a more human-centered approach should focus on measuring (and perhaps, minimizing) users' affective discomfort, rather than automating and approximating their prior decisions in which this affective discomfort is habituated. Work in psychology has found affect—i.e., how a user *feels* about a decision—to be highly predictive of the decisions themselves [20]. In our work, we challenge existing models of whether users are simply surprised by or would accept a permission, and, through an experience sampling method, hope to understand how users *feel* about PUDs in contexts outside of their initial or direct interactions with permissions.

III. Methodology

We conducted a four-week long experience sampling study (ESM)—similar to methods employed in prior work on permission decisions [12]—with 44 people who used an Android smartphone as their primary phone. We deployed an Android application as a study tool: twice per day at random times, we prompted our participants to fill out a questionnaire inquiring about their affective response to the use of a randomly selected permission by a specific application installed on their device. Concurrently, we asked participants to report their current physical location and what they were currently doing (henceforth referred to as "location" and "activity", respectively, and "physical context" jointly). With this collected data, we statistically modeled participants' affective responses to PUDs as a function of physical context, permission type, and app type.

A. Recruitment

We recruited 44 participants via our institution's research pool after a screening them for study eligibility requirements (i.e., left home multiple times per week, were over 18 years old, lived in the United States). Our participants were all from the United States and over the age of 18. A majority of them (30) were aged 35 to 54 years old. 36 participants identified as men, and 8 as women. About a third (14) used more than one phone in their daily life. Most reported being college-educated (35) and employed (42). A majority (23) self-reported spending over 11 hours away from their workplace or home every week.

B. Experience Sampling Method

Past work capturing user acceptance of permissions has automatically prompted users to answer questions at the time of an app's permission checks [2], [5]. However, with this approach, context is tethered to whenever permissions are checked, rather than being sampled randomly. To better understand the correlation between physical context and affective discomfort with PUDs, we needed to sample context randomly, which we achieved by employing the Experience Sampling Method (ESM). This method consists of asking users to self-report on their experiences at random times of day without them expecting it, in the hopes of capturing candid, in-situ responses [12]. In our study, we set up our Android ESM app to randomly prompt users to answer questionnaires throughout the day.

While prior work collecting contextual data via automated questionnaires and telemetry [5], [6] often have sample sizes in the thousands, the contexts we sought to capture—semantic location and activity—and the dependent variable we sought to measure—affective response—would not be possible to automatically infer without significant accuracy and/or privacy concerns. Accordingly, our questionnaires needed to actively disrupt the user experience to collect self-reported data. We limited our prompts to two per day to reduce participant burden.

C. Procedure

Participants were asked to install our ESM Android application on their personal smartphone. Our study procedure was carried out via this application, which had two simultaneous functions: (1) to gather a list of all the installed applications on participants' devices and the permissions they request; and (2) to prompt participants to fill out a questionnaire two times per day.

1) Application data gathered: The Android application we developed collected the name of apps installed on the phone, the permissions the apps requested, and ecological momentary assessments (i.e., self-reported responses to a questionnaire). Over the 30-day participation period, participants were notified to answer two questionnaires per day at random intervals. Questionnaire responses reflected participants' attitudes toward permission requests in *physical* contexts—where or what the participant was doing at the time. Specifically, we collected:

- List of installed applications on participants' phones
- List of permissions the applications request
- Participant responses to our questionnaire, including semantic labels of their location and activity

2) Questionnaire: The twice-daily questionnaire presented to participants comprised of six questions, where we collected a participant's levels of concern and attitude toward a random permission and app combination (which we refer to as a "request"), as well as their self-reported context.

In this work, we analyze only the first and last sections of the questionnaire. In the first section, we asked participants to rate their level of concern with seeing a randomly selected permission request at that moment on a 5-point Likert scale, and then select an affective response, adapted from the PANAS-X [21], that best reflected their attitude toward the permission request (see Q2 in Appendix A for specifics of these responses). In the last section, we ask participants to self-report their physical location and activity at the time of answering the questionnaire. We provide a set of predetermined locations and activities to select (Table I) along with a free-response "other" option. [2]

The full questionnaire is provided in the Appendix A.

3) Ethics and compensation: Our study protocol was reviewed and approved by an institutional review board (IRB). Participants could earn up to $125 in total compensation by the end of the 4-week period. To encourage longitudinal participation, we awarded participants $30 for every 10 days of participation; if they completed the entire study, they would receive a $35 bonus.

IV. Results

In this section, we answer our main RQ: **How does a user's affective discomfort with permissioned uses of data vary across physical contexts?**

We first wanted to understand how location and activity conditions, as well as the interactions therein, could be associated with different levels of user concern with PUDs. To address the former, we generated clusters of participants based on their distributions of locations and activities and compared the clusters' mean concern scores. To address the latter, we constructed an explanatory mixed effects model. These analyses correspond with our aforementioned hypotheses:

H1.1 Participants with different compositions of recorded contexts will have different average levels of concern.

H1.2 Permission requests made in different physical contexts will be correlated with different levels of concern.

A. Understanding concern scores

Before we can evaluate how concern levels might vary across contexts, we must first understand how to interpret the concern scores themselves. Thus, we analyzed each of the five concern score levels by inspecting the different qualitative attitudes that participants self-reported in tandem with their level of concern. The distributions for each level of concern can be found in Figure 1.

[2] The set of activities was inspired by those detectable from Google's Android API, but we did not use the API to fill the form. Participants self-reported their responses. (i.e. DetectedActivity https://developers.google.com/android/reference/com/google/android/gms/location/DetectedActivity).

Locations	Activities	App Types	Permissions
Home	Working	Media	Communications
Car	Commuting	Business/Productivity	Accounts
Friends	Errands	Communication	Activity Sensing
Public Transit	Playing a game	Education	Call Phone
Restaurant	Resting	Finance	Camera
Store	Socializing	Food/Drink	Data Read
Street	Watching TV	Games	Data Write
Work	Other	Health/Fitness	Location
Other		Lifestyle	Receive Messages
		Photo/Video	Record Audio
		Shopping	System Alert Window
		Social	Wake Lock
		Tools	Other
		Travel/Maps	
		Other	

TABLE I

Categories of data collected. Locations and activities were self-reported by participants; app type and permissions were randomly generated for each ESM questionnaire.

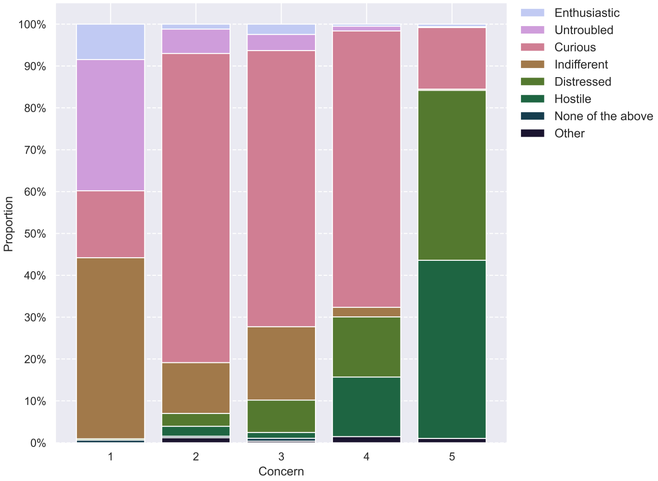

Fig. 1. Distributions of self-reported attitude across all participants responses, grouped by level of concern. Concern levels of 2 and 3 appear relatively similar, with a deviation occurring at level 4, where proportion of "distressed" and "hostile" attitudes increase noticeably, and "indifferent" attitudes decrease.

We found that at concern = 1, participants largely reported being untroubled or indifferent to a given permission request. The attitude compositions at concern levels 2 and 3 were fairly similar to each other, with greater proportions of curiosity, and much less indifference and untroubledness. There was then a marked increase in negative attitudes at concern = 4 and above: participants selected emotions such as hostility and distress at far greater percentages than at lower levels of concern. Owing to the marked discrepancy in attitudinal response between levels 3 and 4, we concluded that a self-reported concern of 1,2, or 3 reflected low affective discomfort, whereas concern levels of 4 or 5 reflected high affective discomfort.

B. *Typology of participants (H1.1)*

To first understand if different compositions of location and activity conditions could be associated with varying levels of user concern, we clustered participants based on their distributions of recorded locations and activities with a K-means clustering algorithm with three clusters (i.e., K=3). For example, participants who recorded a large percentage of their survey submissions at home and less so at work might be clustered with others who split their locations similarly. We chose three clusters based on prior literature on creating permission profiles for users [19], which found that as few as three clusters could be optimal for identifying interpretable clusters. After we clustered participants, we explored whether average levels of concern differed between clusters. We found support for **H1.1**, with average levels of concern varying across context clusters.

Participant profiles based on location distributions mainly varied by the amounts of diary entries recorded at home versus work. The differences in these distributions were also associated with different levels of concern about PUDs:

- **L1**: Majority at home (56%), smaller amounts at work (24%) and in the car (6%). Lowest mean levels of concern (3.19).
- **L2**: Almost entirely at home (87% of responses on average), very little elsewhere. Medium mean levels of concern (3.44).
- **L3**: More even split between home (32%) and work (44%), with small amounts at friends' homes, restaurants, stores, and public transit. Highest levels of concern (mean = 4.08).

Somewhat similarly, participant clusters based on activity distributions varied highly based on diary entries recorded while working versus resting:

- **A1**: Very little working activity (14%); majority resting, watching TV, or "other" activities (77% in sum)—upon inspecting these free text "other" responses, we found that a large majority of them were related to eating or cooking. Lowest mean levels of concern (2.80).
- **A2**: Some working (28%), and similar amounts resting (34%), plus smaller amounts of commuting and errands

(12% in sum). Second-highest mean levels of concern (3.53).

- **A3**: Almost half working (45%); rest of activities evenly split. Highest mean levels of concern (3.59).

While we did not collect additional qualitative data about participants' personal lives or make associations with their demographic characteristics, these clusters can offer a glimpse into the types of lifestyles they might lead. For example, L3, with higher amounts of activity on public transit, at restaurants, and at friends' homes, might point to an urban professional lifestyle. A1, with majority leisure activities and very little working activity, might signal unemployment or being a student. Future work could explore how qualitative differences in lifestyles might be associated with different levels of overall concern with PUDs.

Category	**Variable**	**Coefficient**
Constant	*Home, Working, Media, Comms*	2.930
Location	Car	-0.100
	Friends	0.054
	Public Transit	0.248
	Restaurant	0.289
	Store	0.493*
	Street	-0.354
	Work	0.113
	Other	-0.337*
Activity	Commuting	0.186
	Errands	-0.244*
	Exercising	0.163
	Playing a game	0.368*
	Resting	-0.088
	Socializing	-0.186
	Watching TV	-0.030
	Other	-0.015
App Type	Business & Productivity	-0.095
	Communication	-0.058
	Education	0.166
	Finance	0.428*
	Food & Drink	0.520*
	Games	0.476*
	Health & Fitness	0.111
	Lifestyle	0.204
	Photo & Video	0.462*
	Shopping	0.504*
	Social	0.485*
	Tools	0.184*
	Travel, Maps, Navigation	0.173
	Other or Unknown	0.162
Permission Type	Accounts	0.544*
	Activity Sensing	-0.102
	Call Phone	1.508*
	Camera	1.029*
	Data Read	0.536*
	Data Write	0.324*
	Location	-0.351*
	Receive Messages	-0.341
	Record Audio	1.169*
	System Alert Window	0.903*
	Wake Lock	0.454*
	Group Var	0.440*

TABLE II
Results of the explanatory model (*indicates $p < .05$). The variance in random intercepts was 0.3978. We used the modal category in each variable or covariate—home (location), working (activity), media (app type), connectivity/communications (permission type)—as a baseline.

C. *Effects of specific contexts (H1.2)*

To further understand how specific combinations of physical contextual factors affected user concern, we constructed a mixed effects model, where we included a random intercepts term due to repeated observations from each of our participants. Our independent variables were location and activity, with app type and permission type as covariates. We converted these variables from categorical to binary, with the most common category in each variable or covariate—home (location), working (activity), media (app type), connectivity/communications (permission type)—used as the baseline reference level. In other words, the regression coefficients should be viewed in comparison to a participant at home, working, using a media app that requests a connectivity-related permission. For location and activity, participants also had the option to input a free text "other". For app types and permissions, "other" refers to values with very low counts that were collapsed into one category. The dependent variable was the participant-reported level of concern for a given combination of the independent variables and covariates.

Our results let us understand which permission, location, and activity tuples are correlated with the highest concern for users. A complete view of all the coefficients of our mixed linear model is in Table II. The coefficients represent changes in the level of user concern for a PUD: a positive coefficient means that variable is associated with an increase in concern, and a negative coefficient is associated with a decrease in concern. A first step in comparing effects of each of these coefficients is interpreting the concern associated with the baseline participant permission experience—located at home, working, and being asked for a connectivity-related permission for a media app. The baseline experience was associated with a mean concern score of approximately 2.9. Based on our previous analysis in Figure 1, this score indicates that the baseline experience would most likely be correlated with neutral reactions from a participant: neither particularly high nor low affective discomfort.

In comparison, permission requests of different types and initiated by various app types were associated with significantly higher concern; for example, the "Call Phone" permission was associated with an average concern score of over 1.5 above the baseline (p¡0.001). That multiple permission types have higher concern scores than connectivity-related ones follows some intuition. A mobile device often has Internet, Bluetooth, and NFC enabled simultaneously and by default, so users might not feel particularly concerned about permissioned use of these resources; on the other hand, other permissions might be tied to expectations of specific use cases, making requests for them more jarring. To paraphrase an example from prior work [14], users might feel unconcerned when playing a blackjack game that asks to connect to the Internet, but more alarmed if the same blackjack app is asking to use the camera or make a phone call.

Mismatches between expectations of an app and the permissions it requests can also translate to physical context.

When participants were playing games, their concern levels were also higher (0.368 higher than the baseline). This mirrors outcomes from when permission requests specifically came from gaming-related apps (0.476 higher than the baseline). Here, both instances of heightened concern might be explained by users having expectations of playing games that should not involve any PUDs: to extend the previous example, users might not understand why their blackjack game app needs to know their location when they are out on running errands, and may thus feel heightened concerned over that PUD.

The effects of other contextual variables were more diverse. Permission requests when participants were at a store, for example, had a mean concern score almost 0.5 greater than the baseline. At the same time, running errands was associated with decreasing concern scores by 0.24. Free-response contextual factors (i.e., "Other") were also associated with significantly lower levels of concern (0.337 lower than the baseline). A majority of these responses were related to travel in some way (e.g., hotel, airport); the negative correlation suggests lower concern when in locations habitually distinct for the participant. When participants are in a new environment, their minds are likely more occupied with information about their surroundings rather than permission requests from their mobile device, so they might feel less privacy concern [22].

In other words, knowing a combination of app type and permission type on their own might alert us to use cases that are particularly concerning to a user, and parallel environmental factors might help confirm these heightened concern scenarios. However, other specific contextual factors might point to situations that are *less* concerning, where a user might be willing to allow—or at least be less disquieted by—those same permissions and apps.

V. Discussion

To summarize our findings, we collected and analyzed in-situ affective responses to permissioned uses of data to understand whether and how user concern over these uses varies across location, activity, and permissions. We found distinguishable differences in concern across public versus private physical environments, as well as heightened levels of concern when a permission request did not match any obvious expectations of a physical context. Here, we examine the role of "concern" in user permission preferences, as well as ramifications of dynamic, context-aware applications of affect in user permissions.

A. Predictive interfaces based solely on prior decisions risk automating and normalizing affective discomfort

A natural takeaway from our findings is that we should make permissions more context-aware, but doing so adds additional complexity and decision fatigue to permission interfaces. Accordingly, a fair amount of prior work has explored easing user burden by automating permission configurations based on a users' prior decisions in context (e.g., [3]) or based on predicting the user's preference similarity to other users (e.g., [23]). These predictive interfaces are undeniably useful at reducing immediate user burden, but may pose more sinister long term risks by automating and normalizing affective discomfort. Indeed, as we note in Section II, just because a user *accepts* a permission doesn't mean they are not concerned about it. For example, users might grant permissions under financial duress [8].

Continuing to focus on automation based solely on prior acceptance of permissions or based on other users' decisions could further encode this mis-alignment, causing a feedback loop where users end up normalizing what they find creepy just because they have accepted it in the past. This normalization of affective discomfort has sinister externalities: users' intent to use privacy-invasive apps aligns with what they think is "normal" in society, rather than their genuine discomfort with the invasiveness [7]. As researchers begin to explore the potential for intelligent assistants on mobile devices to become more proactive in protecting users' privacy [24], [25] and the reach of mobile permissions expands into augmented reality environments [26], accounting for affective discomfort when automating decision making may be more necessary than ever.

B. Measuring affective discomfort to help users make better, context-specific permission decisions

Using concern and/or affective discomfort as a proxy for user permission preferences fits into the broader field of affective computing [27], the study of computing systems that specifically interact with and interpret a user's internal state, or affects. In this genre of HCI, interfaces respond to a user's affect by changing affect, adapting to affect, generating affective behavior by the machine, or modeling the user's affective state. As we hope to design more user-centric interactions and systems, it is only natural that we take the user's affect into account [28].

How a user *feels* about a decision can sometimes be used to predict the decision itself [20]. However, concern is, ostensibly, only one part of how a user feels about allowing or denying permissions Even if there is low concern around a permission request context, a user might still not necessarily wish to give it access anyway. Thus we are not advocating for using *only* measurements of concern when aiming to facilitate user decision making; rather, we argue that without taking concern into account and focusing only on automation based on prior decision making, we risk perpetuating affective discomfort in mobile app use. Thus, further highlighting and measuring this discomfort, as well as taking steps to *reduce* it, are worthy endeavors in the future.

Measurements of concern also open new design opportunities for educational interventions. Empirical studies have shown that users are generally ill-informed about what permissions do [10], poising them to make decisions that might be against their best interest, e.g., accepting all permissions. When we asked our participants to provide an emotional response to a PUD, many who felt concerned also indicated curiosity about the PUD. This result suggests that users might prefer to be better informed about how and why PUDs occur, rather than simply presented with the decision or knowledge that a PUD is happening. Dynamic approaches to permissions could inform users when and why changes to permissions are made based on changes in context, rather than being a black box. By pairing a personalization engine with education, users can make better decisions that feedback into the system.

VI. Limitations

A few limitations arose in our work. In our questionnaire, we did not limit the set of permissions we asked participants about only to those that the participant had already allowed. As such, certain PUDs might feel particularly alarming if the participant remembered explicitly denying those permissions in the past. However, we still were able to study the main effect we hypothesized: that there would be differences in user concern over PUDs across different physical contexts. Also, while our participants each contributed a large quantity of questionnaire responses, we could not obtain representative samples for every single combination of context, app type, and permission available. Relatedly, our participants were overwhelmingly male and college-educated; more diverse sampling could help make more generalizable claims about differences in concerning contexts.

VII. Conclusion

Through a four-week long experience sampling study, we studied how users' concerns over permissioned uses of data (PUDs) by third-party smartphone applications varied across physical contexts and environments. We found distinguishable differences in concern for PUDS in public versus private contexts. We also found that participants with similar contextual compositions in their day-to-day lives reported similar overall levels of concern with PUDs. These findings suggest that permission decisions made in one context should not necessarily translate to others; rather, context-aware permission policies may help reduce user concern. We discussed the ramifications of such an approach, recommending that care be taken to integrate user agency when designing automated permission systems. In short, our work lays the foundation for future work exploring human-centered context-aware permission controls to reduce affective discomfort and concern with permissions.

Acknowledgment

This work was generously funded, in part, by NSF SaTC Grants #2316294 and #2316287.

References

[1] L. Tsai, P. Wijesekera, J. Reardon, I. Reyes, S. Egelman, D. Wagner, N. Good, and J.-W. Chen, "Turtle guard: Helping android users apply contextual privacy preferences," in *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, 2017, pp. 145–162.

[2] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov, "The feasibility of dynamically granted permissions: Aligning mobile privacy with user preferences," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 1077–1093.

[3] P. Wijesekera, J. Reardon, I. Reyes, L. Tsai, J.-W. Chen, N. Good, D. Wagner, K. Beznosov, and S. Egelman, "Contextualizing privacy decisions for better prediction (and protection)," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–13.

[4] E. Hayashi, S. Das, S. Amini, J. Hong, and I. Oakley, "Casa: context-aware scalable authentication," in *Proceedings of the Ninth Symposium on Usable Privacy and Security*, 2013, pp. 1–10.

[5] R. Mendes, A. Brandão, J. P. Vilela, and A. R. Beresford, "Effect of user expectation on mobile app privacy: a field study," in *2022 IEEE international conference on pervasive computing and communications (PerCom)*. IEEE, 2022, pp. 207–214.

[6] R. Mendes, M. Cunha, J. P. Vilela, and A. R. Beresford, "Enhancing user privacy in mobile devices through prediction of privacy preferences," in *Computer Security–ESORICS 2022: 27th European Symposium on Research in Computer Security, Copenhagen, Denmark, September 26–30, 2022, Proceedings, Part I*. Springer, 2022, pp. 153–172.

[7] J. S. Seberger, I. Shklovski, E. Swiatek, and S. Patil, "Still creepy after all these years: The normalization of affective discomfort in app use," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–19.

[8] C. W. Munyendo, Y. Acar, and A. J. Aviv, ""desperate times call for desperate measures": User concerns with mobile loan apps in kenya," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2022, pp. 1521–1521.

[9] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall, "A conundrum of permissions: installing applications on an android smartphone," in *International conference on financial cryptography and data security*. Springer, 2012, pp. 68–79.

[10] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner, "Android permissions: User attention, comprehension, and behavior," in *Proceedings of the eighth symposium on usable privacy and security*, 2012, pp. 1–14.

[11] P. G. Kelley, L. F. Cranor, and N. Sadeh, "Privacy as part of the app decision-making process," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2013, pp. 3393–3402.

[12] B. Bonné, S. T. Peddinti, I. Bilogrevic, and N. Taft, "Exploring decision making with {Android's} runtime permission dialogs using in-context surveys," 2017, pp. 195–210. [Online]. Available: https://www.usenix.org/conference/soups2017/technical-sessions/presentation/bonne

[13] W. Cao, C. Xia, S. T. Peddinti, D. Lie, N. Taft, and L. M. Austin, "A Large Scale Study of User Behavior, Expectations and Engagement with Android Permissions," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 803–820.

[14] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing," in *Proceedings of the 2012 ACM conference on ubiquitous computing*, 2012, pp. 501–510.

[15] A. Gorla, I. Tavecchia, F. Gross, and A. Zeller, "Checking app behavior against app descriptions," in *Proceedings of the 36th international conference on software engineering*, 2014, pp. 1025–1035.

[16] A. Rao, F. Schaub, N. Sadeh, A. Acquisti, and R. Kang, "Expecting the unexpected: Understanding mismatched privacy expectations online," in *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 2016, pp. 77–96.

[17] P. H. Chia, Y. Yamamoto, and N. Asokan, "Is this app safe? a large scale study on application permissions and risk signals," in *Proceedings of the 21st international conference on World Wide Web*, 2012, pp. 311–320.

[18] A. Alsoubai, R. Ghaiumy Anaraky, Y. Li, X. Page, B. Knijnenburg, and P. J. Wisniewski, "Permission vs. app limiters: Profiling smartphone users to understand differing strategies for mobile privacy management," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–18.

[19] B. Liu, J. Lin, and N. Sadeh, "Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help?" in *Proceedings of the 23rd international conference on World wide web*, 2014, pp. 201–212.

[20] C. J. Charpentier, J.-E. De Neve, X. Li, J. P. Roiser, and T. Sharot, "Models of Affective Decision Making: How Do Feelings Predict Choice?" *Psychological Science*, vol. 27, no. 6, pp. 763–775, Jun. 2016, publisher: SAGE Publications Inc. [Online]. Available: https://doi.org/10.1177/0956797616634654

[21] D. Watson and L. A. Clark, "The panas-x: Manual for the positive and negative affect schedule-expanded form," 1994.

[22] A. Ioannou, I. Tussyadiah, and Y. Lu, "Privacy concerns and disclosure of biometric and behavioral data for travel," *International Journal of Information Management*, vol. 54, p. 102122, 2020.

[23] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow my recommendations:

A personalized privacy assistant for mobile app permissions," in *Twelfth symposium on usable privacy and security (SOUPS 2016)*, 2016, pp. 27–41.

[24] N. Malkin, D. Wagner, and S. Egelman, "Runtime permissions for privacy in proactive intelligent assistants," in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022, pp. 633–651.

[25] A. Stöver, S. Hahn, F. Kretschmer, and N. Gerber, "Investigating how users imagine their personal privacy assistant," *Proc. Priv. Enhancing Technol*, vol. 2, pp. 384–402, 2023.

[26] D. Harborth and A. Frik, "Evaluating and redefining smartphone permissions with contextualized justifications for mobile augmented reality apps," in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021, pp. 513–534.

[27] R. W. Picard, *Affective computing*. MIT press, 2000.

[28] E. Hudlicka, "To feel or not to feel: The role of affect in human–computer interaction," *International Journal of Human-Computer Studies*, vol. 59, no. 1, pp. 1–32, Jul. 2003. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1071581903000478

Appendix

Q1 Imagine that {APPLICATION NAME} is using {PERMISSION REQUEST CANONICALIZATION} right at this moment. On a scale from 1 - 5, where 1 represents "not at all concerned" and 5 represents "very concerned", how concerned would you be about this access? *(1 - 5 Likert scale)*

Q2 Among the list below, please check off words that describe your attitude towards this access. Select all that apply.

- [ ] Enthusiastic,
- [ ] Distressed,
- [ ] Curious,
- [ ] Indifferent,
- [ ] Hostile,
- [ ] Untroubled,
- [ ] None of the above,
- [ ] Other

Q3 ***EXCLUDED FROM ANALYSIS*** Now imagine that the application, {NAME}, is {PERMISSION REQUEST CANONICALIZATION} in order to {MALICIOUS/BENIGN PURPOSE}. On a scale from 1 - 5, where 1 represents "not at all concerned" and 5 represents "very concerned", how concerned would you be about this access? *(1 - 5 Likert scale)*

Q4 ***EXCLUDED FROM ANALYSIS*** Which of the following restrictions would *reduce* your concern with this access? Please select all that apply.

- [ ] Only if I am actively using the app
- [ ] Only while I'm at a pre-specified location (e.g., work, home)
- [ ] Only at a certain time of day (e.g., between 9am and 5pm)
- [ ] Only on certain days of the week (e.g., only on weekends)
- [ ] Only while I'm doing a certain activity (e.g., running, browsing the web)
- [ ] I am not concerned by this access in any context
- [ ] This app should never be able to have this access
- [ ] Other

Q5 In a word or two, please describe your current location (e.g., home, work, coffee shop, gym).

- [ ] Home
- [ ] Work
- [ ] Gym
- [ ] Friend's
- [ ] Restaurant
- [ ] Store
- [ ] Car
- [ ] Public Transit
- [ ] Street
- [ ] Other

Q6 In a word or two, please describe what you are currently doing (e.g., "working", "exercising", "socializing").

- [ ] Commuting
- [ ] Errands
- [ ] Socializing
- [ ] Working
- [ ] Playing a game
- [ ] Resting
- [ ] Watching TV
- [ ] Exercising
- [ ] Other