Contextual Permission Models for Better Privacy Protection

Primal Wijesekera[1,2], Joel Reardon[3], Irwin Reyes[2], Lynn Tsai[1], Jung-Wei Chen[4],
Nathan Good[4], David Wagner[1], Konstantin Beznosov[5], and Serge Egelman[1,2]
[1]University of California, Berkeley, CA [2]International Computer Science Institute, Berkeley, CA
[3]University of Calgary, Calgary, AB [4]Good Research, Berkeley, CA
[5]University of British Columbia, Vancouver, BC
{primal,lynntsai}@berkeley.edu, joel.reardon@ucalgary.ca, ioreyes@icsi.berkeley.edu,
{jennifer,nathan}@goodresearch.com, {daw,egelman}@cs.berkeley.edu, beznosov@ece.ubc.ca

ABSTRACT

Mobile operating systems have permission systems that regulate how apps access sensitive device resources and user data. Modern permission systems, however, are not always well-aligned with users' privacy expectations: users are often unaware of how often and under what circumstances apps access sensitive data on the device, and may have privacy preferences that vary under different circumstances. We developed a method to systematically reduce this disconnect between user privacy expectations and reality. In evaluating our methods, we found that a significant portion of users make *contextual* privacy decisions: when determining the appropriateness of an app accessing sensitive data, users consider what they were doing on their phones at the time, such as whether or not they were actively using the app requesting the data. Existing privacy mechanisms fail to account for these contextual factors, resulting in privacy violations, as user expectations are not being met. Our work provides further empirical evidence that Contextual Integrity is a viable framework for minimizing privacy violations on mobile devices, but that future work is needed to systematically show how it can be implemented in a wide range of systems. In attaining that goal, we show that machine learning can be used in permission systems to help account for context, reducing privacy violations by 80% and the need for user involvement.

CCS CONCEPTS

• **Security and privacy → Information flow control**; **Usability in security and privacy**; • **Human-centered computing → Empirical studies in ubiquitous and mobile computing**;

KEYWORDS

Privacy, mobile permissions, access control, user study

1 INTRODUCTION

While large corporate data breaches and intrusions garner much media attention, ordinary end-users are also targets of complex threats to their security and privacy. Smartphones have inadvertently become a major enabler of privacy breaches. Mobile platforms have developed permission systems that empower the user to regulate access to sensitive data and resources on the device. Modern permission systems, however, lack the ability to take into account users' privacy expectations under different contexts, resulting in a significant gap between how these permission systems perform and how users would like the platform to protect their sensitive data. This work offers a summary of our prior research on how users make privacy decisions on smartphones and how the mobile platforms can better meet users' privacy requirements by applying the theory of Contextual Integrity (CI) [2].

2 QUANTIFYING USER EXPECTATIONS

In order to be able to infer contextual information to better protect sensitive resources, it is imperative to understand the different contexts under which sensitive resources are accessed in the wild, and how user expectations change as a function of these different contexts. We performed a 36-participant field study to quantify how often mobile apps access protected resources when users are not expecting it [3]. Participants used specially-instrumented Android phones that logged every time a sensitive resource was accessed and some surrounding contextual information. In exit interviews, we asked participants to review a selection of screenshots taken when an app accessed a given sensitive resource and to state whether they would have allowed (or denied) that access had they been given the chance. Over 80% of participants responded that they would have denied at least one such resource request. We also found that users were more likely to deny access in circumstances when the requesting app was running invisibly in the background (i.e., when the participant was likely unaware that the requesting app was even running). At the time of the study, Android had "ask-on-install" permissions, where users were forced to make a decision during app installation with little or no information about how and when the app would access sensitive data in the future.

To the best of our knowledge, this was the first field study of how apps access sensitive resources and whether user expectations align with actual app behaviors. The biggest contribution of the work, in terms fo the CI framework, was to quantitatively show that the visibility of the requesting app has a significant impact on users' expectations. On average, apps requested access to permission-protected resources over 200 times per hour. Thus, it is impractical to burden the user with runtime prompts for each permission request; satisfying users' desires for more contextual control over app permissions requires a different approach besides asking the user to manually evaluate every permission as apps request it (or only doing so once, during app installation, in the absence of more meaningful contextual information).

3 PREDICTING PREFERENCES

In measuring real-world app behaviors and user expectations, we found that users are likely to take contextual factors into account when deciding whether to allow or deny a request. However, it is

impractical to prompt the user to make these decisions every single time. One plausible alternative would be to involve the user in the beginning and use their initial responses to predict their future expectations under varying situations.

We performed a 131-participant field study to explore the feasibility of predicting users' future privacy decisions based on their past decisions and available contextual information [4]. Just as in the previous work, participants used specially-instrumented phones that logged whenever an application requested access to sensitive resources. We used the Experience Sampling Method (ESM) to collect ground truth data about users' privacy preferences [1]. We probabilistically asked them about an app's recent access to data on their phone, and whether they would have permitted it if given the choice. We treated participants' responses to these ESM probes as the main dependent variable in our machine learning model. Through this study we collected 175M data points including 4K+ real-world user decisions.

We used our field study data to build a classifier that attempts to predict users' privacy decisions based on past responses to prompts, the data requested, and certain contextual information surrounding the request (i.e., the visibility of the requesting app and the foreground application at the time of the request, if different). We demonstrated that this classifier-based approach can correctly predict users' privacy decisions 96.8% of the time. This produces an 80% reduction in mismatches between app behaviors and user preferences, as compared to the "ask-on-first-user" (AOFU) approach,[1] which is the current standard on Android and iOS. Our ML-based approach also reduced user involvement by 20%.

Any attempt to understand how context affects users' decision-making needs to first identify the contextual factors that help users to make varying decisions. Our prior studies identified both the visibility of the requesting app and the foreground app (if different) [2] at the time of the request as having the highest information gain in the proposed machine learning model. The significance of this observation is that future permissions systems will likely need to present this information to users in a meaningful way, so that they can make more informed decisions. Yet, much more research is needed to determine the other contextual factors that can be incorporated into similar classifiers.

4 CONTEXTUAL PERMISSION SYSTEMS

Based on our findings, we developed a custom Android version with a contextually-aware permission model [5]. In our previous work, the proposed machine learning model was an offline classifier trained on our collected data. In this study, we implemented the entire machine learning pipeline in Android so that the platform could make real-time decisions on sensitive permission requests. The new model guards resources based on the user's past decisions under similar circumstances, as inferred from the collected contextual information. We performed a 38-person field study to measure the efficiency and usability of the new permission model. Based on exit interviews and 5M data points, we found that the new system is effective in reducing potential privacy violations by 75%. Despite being significantly more restrictive over the default permission system, participants did not find that the new model caused any usability issues. That is, automatically denying apps access to data at runtime did not result in unexpected app behavior, such as crashes or loss in functionality. This suggests that the classifier was correctly denying apps access to only data that was not absolutely necessary.

To the best of our knowledge, we are the first to implement a real-world contextually aware permission model that takes surrounding contextual signals into account before deciding on allowing (or denying) a given permission request. We show that if suitable data spoofing [3] mechanisms are in-place, a restrictive permission model can preserve both user privacy and app functionality.

We are interested in participating in this workshop both to present this relevant prior work, as well as to solicit ideas for next steps, so that we can better apply the theoretical framework to our system architecture.

REFERENCES

[1] Stefan E Hormuth. 1986. The sampling of experiences in situ. *Journal of personality* 54, 1 (1986), 262–293.

[2] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (February 2004), 119.

[3] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android Permissions Remystified: A Field Study on Contextual Integrity.. In *USENIX Security Symposium*. 499–514.

[4] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov. 2017. The Feasibility of Dynamically Granted Permissions: Aligning Mobile Privacy with User Preferences. In *2017 IEEE Symposium on Security and Privacy (SP)*. 1077–1093. https://doi.org/10.1109/SP.2017.51

[5] Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman. 2018. Contextualizing Privacy Decisions for Better Prediction (and Protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 268, 13 pages. https://doi.org/10.1145/3173574.3173842

[1] AOFU prompts the user whenever an app accesses a resource for the *first* time, re-applying the user's decision to all subsequent requests.

[2] The "foreground app" is the app the user is engaged with at the time of the request, which may be different than the app initiating the request.

[3] Rather than outright denying access to data, our system sometimes chooses to inject fake data into apps.