





Visual Attention and Privacy Indicators in Android: Insights from Eye Tracking

Michele Guerra¹, Roberto Milanese^{1,2}, Michele Deodato³, Vittorio Perozzi¹
and Fausto Fasano¹

¹*Mosaic Research Center, University of Molise, Italy*

²*Department of Control and Computer Engineering, Politecnico di Torino, Italy*

³*Division of Science, New York University, Abu Dhabi, U.S.A.*

Keywords: Security and Privacy, Application Security, App Permission, Android, Android Permission Model.

Abstract: In today's digital landscape, where privacy preservation is of paramount importance, Android has implemented new features to enhance transparency: the Privacy Indicators (PIs). Our study employs eye-tracking technology to investigate how users perceive and interact with these indicators. As a visual alert system, PIs signal when sensitive resources, like camera or microphone, are in use. However, the structure of Android's permission model, susceptible to exploitation by malevolent or commercial apps, places an excessive responsibility on PIs. They act as the final alert for users against the misuse of permissions in unexpected contexts. We conducted a controlled experiment with 29 participants who were exposed to various privacy scenarios while their eye movements were tracked and recorded. Our findings reveal a significant gap in PIs effectiveness, particularly in high-engagement tasks, indicating a need for more eye-catching privacy notifications. These findings suggest the need for redesigning some privacy interfaces to make them more effective. The study's insights contribute to the broader discussion on balancing functionality with user privacy and the methodology of utilizing eye tracking in user experience research.

1 INTRODUCTION

Android, the undisputed leader in the mobile device world, has become an integral part of our everyday lives, with its ubiquitous operating system deeply entrenched into the fabric of our digital experiences. The platform's widespread adoption brings to the fore critical questions about user privacy and the efficacy of mechanisms designed to protect it. Android's permission model, a framework established to regulate app access to sensitive resources like cameras and microphones, is central to this discussion. However, users' growing trust in applications increasingly challenges this system's effectiveness, often leading to a casual approach to granting permissions.


The inherent complexities of Android's permission model and user behavior underscore the need for robust privacy safeguards. This is where Privacy


Indicators (PIs) come into play. Implemented as visual cues within the operating system, PIs aim to alert users when sensitive resources such as camera and microphone are accessed (Figure 1). Yet, questions linger about their actual impact and visibility in everyday use, especially when permissions are exploited in non-contextualized or malicious ways.


PIs became a part of the Android operating system starting from Android 12. They are a critical update in Android's ongoing evolution to enhance user awareness and control over privacy. This study seeks to evaluate the practical impact of this significant feature in real-world scenarios, providing insights into its effectiveness and areas for potential improvement.


Our study seeks to illuminate this crucial aspect by employing eye-tracking technology, a method traditionally reserved for fields such as marketing and medical research, now repurposed to scrutinize user interactions with PIs. This innovative approach allows for an objective, real-time analysis of how users notice and process PIs amidst various interactive tasks on their smartphones.

In this context, we conducted a controlled experi-

^a <https://orcid.org/0009-0005-9990-234X>

^b <https://orcid.org/0009-0009-8758-753X>

^c <https://orcid.org/0000-0002-2624-1430>

^d <https://orcid.org/0000-0003-3736-6383>

ment involving 29 participants to empirically investigate user interaction with Privacy Indicators on Android devices. These participants, drawn from diverse backgrounds, interacted with a series of tasks designed to simulate real-world smartphone usage. Through eye-tracking, our objective was to monitor how and when these users noticed the Privacy Indicators during different types of interactions, especially in scenarios where resource use might not be anticipated.

Through a comprehensive experimental setup, we aim to probe the depths of user attention and response to PIs, assessing their prominence and effectiveness in real-world scenarios. This investigation is particularly timely given the increasing sophistication of apps in seeking permissions under benign guises, only to exploit them for less scrupulous purposes, as explored in the research by (Guerra et al., 2023).



Figure 1: An example of Privacy Indicator in action on Android 12.

Our experiment revealed critical insights into user engagement with PIs. The findings suggest that the classic PIs in their current form are often unnoticed, especially in high-engagement tasks, highlighting a potential vulnerability in privacy protection. This underscores the need for more intuitive and noticeable privacy indicators that can effectively alert users to resource access in real-time.

Our research endeavors to contribute significantly to the discourse on mobile privacy by bridging the gap between technological capability and user behavior. The insights gleaned from this study are expected to inform the design of more intuitive and effective privacy features, steering the mobile computing landscape toward a safer, more privacy-conscious future.

2 RELATED WORKS

2.1 Android Permission Model Evolution and Privacy Concerns

In Android, access to sensitive user data (such as contacts in the address book, received messages, or calendar appointments) or critical system resources (such as the camera, microphone, or biometric sensors) by applications is regulated through the mechanism of permissions (Wang et al., 2021).

The inception of Android's permission model was marked by a simplicity that, while facilitating early user experience, left much to be desired in terms of privacy safeguards. From unrestricted resource access in Android 1.0, the model evolved through subsequent releases. Android 3.0 introduced external storage protections, and Android 4.4 required permissions to be declared at installation, with Android 5.0 adding new permissions into the mix, albeit still at install-time (Felt et al., 2012).

However, these iterations did not offer real-time access management, presenting a dilemma: users could either accept all permissions wholesale or abandon app installation (Peruma et al., 2018). The choice led to an all-or-nothing decision that was insufficient for nuanced privacy management. It was observed that the permissions system's lack of contextualization failed to communicate the risks effectively to the users (Wijesekera et al., 2018).

With Android 6.0, Google reformed its permission model to allow for runtime requests, marking a significant pivot towards user-centric privacy controls (Shen et al., 2021). This approach is intended to provide greater transparency by enabling users to grant permissions based on immediate app usage. Thus, in this Android version, the user can reject an access request (in Android 11, the approval can even be limited to a single usage) and revoke previously approved permissions from the system settings. In the runtime model, permissions are granted to the entire application, rather than to specific features or usage contexts, leading to a prevalent misuse of non-contextualized permissions by a majority of applications (Guerra et al., 2023). This approach limits the control over how, when, and why data access is made (Scoccia et al., 2021). A customized, user-centered permission model can contextualize the request for permission use. This approach has been effectively used to identify and prevent malicious applications holding logic bombs (Fasano et al., 2023). An improvement to the permission model is made in Android 12, which introduced Privacy Indicators and a privacy dashboard to assist users in managing their

data more effectively, aiming to rectify previous models' limitations (Shen et al., 2021). In (Guerra et al., 2023), authors conducted a controlled experiment to assess PIs effectiveness. However, our research introduces a novel methodology by incorporating eye tracking to objectively measure user interaction with Privacy Indicators (PIs). Unlike previous studies that relied on user-reported detection of resource usage, our approach captures real-time, visual engagement data to determine whether users notice the PI when resources like the camera or microphone are accessed. This allows for a more granular analysis of user attention and perception. By tracking eye movement, we can discern not just whether users are aware of resource usage, but also the exact moments and context in which these indicators are observed or overlooked, providing richer insights into the design's efficacy.

2.2 Eye Tracking in Mobile Computing Privacy

The application of eye-tracking technology in understanding user interaction with privacy and security features on digital platforms is emerging as a crucial research direction. Studies by Furman et al. (Furman and Theofanos, 2014) and Egelman et al. (Egelman, 2013) have employed eye tracking to evaluate how different information presentations affect user decisions in authentication scenarios. These studies reveal that while additional information may increase reading time, it doesn't necessarily influence decision-making processes such as using Facebook Connect. The pioneering work by (Punde et al., 2017) demonstrated the broad applications of eye tracking across various fields, providing a foundation for its integration into privacy studies. Building upon this, research by (Carter and Luke, 2020) utilized eye tracking to probe the eye-mind connection, revealing the subconscious processes that guide visual attention and, by extension, privacy-related behaviors on digital platforms. Expertise in computer security also plays a significant role in user engagement with security features. Arianezhad et al. (Arianezhad et al., 2013), monitoring eye movements, found that individuals with security knowledge spend more time looking at security indicators. Similarly, Whalen et al. (Whalen and Inkpen, 2005) demonstrated the tendency for users to overlook security indicators without explicit prompting, underscoring the potential for inattentive blindness in digital security contexts. As eye-tracking devices become more accessible and integrated into everyday technology, active eye-tracking applications are being explored for security purposes. Miyamoto et al. (Miyamoto et al., 2014) developed

EyeBit, an eye-tracking system that encourages users to verify the URL before inputting sensitive data to combat phishing. Installing input fields only after confirming the user's gaze on the URL bar is an innovative approach to instilling secure online behaviors.

Despite the established utility of eye tracking, its application to privacy indicators (PIs) on mobile devices remains underexplored. Conventional studies have simulated user interactions with app prototypes and followed up with questionnaires to gauge resource usage awareness post-interaction (Guerra et al., 2023). Our research diverges from these methodologies by employing real-time eye tracking to capture immediate user responses to PIs. This direct measurement of attention allocation offers a nuanced understanding of user engagement with PIs and addresses the potential shortfalls of self-reported data.

The challenges of ensuring user attention to privacy prompts, as investigated by (Anderson et al., 2016), highlight the importance of overcoming habituation. Our study aims to mitigate this through the objective analysis provided by eye tracking, contrasting previous approaches that lacked this real-time evaluative component. By doing so, we align with the recommendations by (Shen et al., 2021) and (Elbitar et al., 2021) for more personalized and contextually appropriate privacy controls, aiming to present a comprehensive overview of how users perceive and process privacy notifications within their digital experiences.

Our research contributes to this dynamic field by investigating the real-time detection of Privacy Indicators using smartphone-embedded eye-tracking technology. We assess not only the attention drawn by these indicators but also their design and position relative to user gaze patterns during various interactive tasks. By doing so, we address the gap in understanding the immediate impact of PIs on user behavior and awareness in mobile computing.

3 EMPIRICAL STUDY DESIGN

The primary *aim* of our study is to understand user perception of Privacy Indicators (PIs) using eye tracking, given their critical role as defined by Google in signaling live resource use. Specifically, we focus on how users detect and react to PIs that denote unauthorized access to sensitive resources such as cameras or microphones by potentially malicious applications. Additionally, the study explores user responsiveness to two innovative PI designs across various interactive contexts. As the sole alert mechanism for sensitive resource usage,

these indicators present a unique opportunity to measure their conspicuousness and the immediacy of user awareness facilitated by them. The research was driven by the following question:

RQ1: *Do the app characteristics influence the effectiveness of Privacy Indicators (PIs)?*

A controlled experiment involving human subjects was designed to address this RQ.

3.1 Context Selection

The context for our study was crafted to mirror real-world usage where Privacy Indicators (PIs) play a pivotal role in alerting users to the live use of sensitive device resources like cameras or microphones. Unlike simulations of app environments, our experiment utilizes a web application with integrates eye-tracking technology to create a controlled setting that closely replicated user interactions with actual device notifications. The subjects comprised a group of 29 university students, selected for their commonality with the general Android user population. The web application was designed as the interactive 'object' of the study, in an Android environment where classic and alternative PIs were triggered unexpectedly. This consisted of two main categories of tasks: low spatio-temporal attention tasks (quizzes) and high spatio-temporal attention tasks (games). Including these diverse tasks allowed us to evaluate the PIs effectiveness across a spectrum of user engagement levels, providing insights into how attention to PIs might vary with user cognitive load. Unlike previous studies that simulated popular applications, we designed a singular application tailored to the experiment's requirements, ensuring a consistent platform for all participants. Before initiating the experimental tasks, participants provided detailed information about their personal devices, including vendor, model, and operating system, to ensure comprehensive data relevance and to prepare the experimental setup adequately. We excluded devices solely running Apple's iOS to maintain focus on Android's user interface, which could influence the study's outcomes due to cross-platform differences. Within our application, we designed the display of *Classic* privacy indicators, as presently incorporated in the latest Android versions, to trigger without user anticipation. This setup allowed us to explore user perceptions in a setting that mimicked potential malicious application behavior. To better understand if an implementation of more visible PIs can affect the achieved results, we also introduced two innovative PI designs—*edge* and *disk*—each exhibited

twice to the user under various task conditions.

3.2 Experimental Procedure

We explored three distinct PIs (Figure 2), each with unique visual characteristics meant to signal the user about the real-time usage of sensitive resources. Here we detail these indicators, which were the focal point of our investigation, describing their design rationale and expected impact on user experience and privacy awareness:

- **Classic PI:** this is the current implementation of PI in the latest versions of Android. This indicator typically appears as a small icon at the top right edge of the screen, generally within the status bar, indicating the use of either the camera or microphone. In our experiment, this PI was designed to emulate this typical behavior closely. It appeared as a predefined icon – a simple camera or microphone symbol – to signal the operation of these resources. The icon was intended to be discreet yet noticeable enough to inform the users without causing significant distraction from their primary tasks on the device.
- **Disk PI:** this is conceptualized as a dynamic visual cue to enhance the visibility of the resource usage notification. It consists of an outer ring that maintains a constant size while the inner circle size changes, creating a pulsating effect. The color of the disk alternates between green and yellow, providing an additional visual signal of the camera or microphone in use. This design intended to make the classic indicator more visible without radically altering its form or position, presenting a potentially more eye-catching and thus effective notification method.
- **Edge PI:** this is an innovative alternative, inspired by the edge lighting feature found in certain Samsung devices. It is represented as a small, colored slider that animates horizontally across the top edge of the screen, from left to right, and fades out as it moves along its path. The animation includes a fading trail to draw the user's attention subtly, aiming to determine whether motion and transient visual changes could more effectively alert users. The choice to confine the edge effect to the top of the screen was a deliberate design constraint, considering the limitations of measuring visibility for indicators that occupy the full screen on a mobile device.

The experimental procedure of our study was a structured investigation into how users perceive and interact with Privacy Indicators (PIs) while engaging




Figure 2: Comparative Display of Classic Privacy Indicators and Proposed Alternatives.

in distinct tasks within a specialized web application enriched with eye-tracking technology. This bespoke platform was designed to record precise user reactions to classic and alternative PIs under conditions that closely mimic real-world application usage.

The study was conducted in a dedicated university classroom designed to minimize distractions and optimize eye-tracking accuracy. Specific measures, such as intense uniform lighting and a monochromatic backdrop, were taken to ensure high-quality data capture and minimize potential error sources. The participants were instructed on maintaining a stable posture and were briefed about the calibration process and the tasks they would undertake during the experiment. The indicators were displayed briefly at the top of the screen while users were engaged in the designed activities, capturing their attention and response patterns using eye-tracking metrics.

At the beginning, participants were familiarized with the experimental protocol. The experimental procedure was delineated as follows (Figure 4):

- **Initiation:** users start with reading the instructions and granting necessary permissions.
- **Calibration:** the task proceeds with the calibration interface, engaging with the system to train the eye tracker.
- **Task Engagement:** users are perform the assigned tasks, during which PIs can be presented.
- **Interaction Logging:** all user interactions with the tasks and PIs are logged alongside eye-tracking data.

- **Data Submission:** upon completing each task, data are submitted to the server.

After giving participants the instructions and ensuring their consciousness about the experiment, they started the calibration phase to adapt the eye-tracking system to their individual visual and interactive patterns. Each participant was asked to follow and touch specific points on their screen, which were highlighted to attract focus (Figure 3). The calibration phase was a recurring step in the experiment, serving as a gateway between tasks to ensure the continued accuracy of the eye-tracking data. The calibration process was meticulously designed to require a minimum of 70% accuracy. This threshold was established as an optimal balance between precision and usability. Indeed, posing an higher threshold would result in frequent recalibrations, disrupting the experiment flow, while lower values would compromise data integrity. In case participants failed to achieve 70% accuracy, the system prompted a recalibration until the required precision level was reached. After the initial calibration phase, continuous drift correction is applied throughout task execution, refining the calibration with each interface touch to progressively enhance the eye-tracking data accuracy.

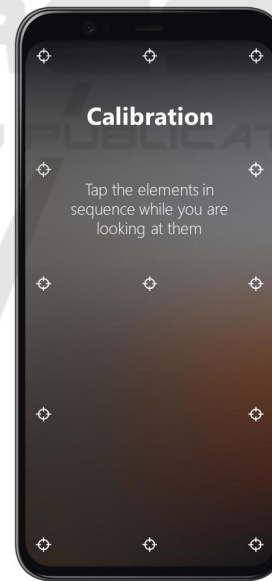


Figure 3: Calibration process screen displaying markers for eye-tracking accuracy improvement.

After the calibration phase, participants engaged in a series of six tasks, alternating between quizzes and games to challenge their attention and cognitive load (Figure 5). These tasks were intended to explore the effectiveness of PI notifications within apps that significantly differ in terms of user engagements:




Figure 4: Sequential Diagram of the Controlled Experiment Process.

low spatio-temporal attention during quizzes and high spatio-temporal attention during interactive gaming scenarios. The sequence of tasks was carefully orchestrated as follows:

- Q_1 [**Classic PI**] Participants began with a quiz designed to assess their focus on content and subsequent noticeability of the *Classic PI*.
- G_1 [**Edge PI**] The next task was an interactive game intended to examine how the alternative *Edge PI* performed in an environment requiring high user interaction.
- Q_2 [**Disk PI**] Another quiz session followed, this time using the *Disk PI*, to test its visibility in a low spatio-temporal attention context.
- G_2 [**Classic PI**] The sequence returned to a gaming task with the *Classic PI* to evaluate consistency across task types.
- Q_3 [**Edge PI**] A repeat of the quiz format provided another data point for the *Edge PI*, reassessing its effectiveness.
- G_3 [**Disk PI**] Finally, the *Disk PI* was revisited during a gaming task to complete the set of experimental conditions.

To simulate a realistic usage scenario, PIs and their alternatives were displayed for exactly eight seconds, randomly materializing within a three seconds interval following the onset of a corresponding task. This randomness introduced an element of unpredictability, mimicking real-world notifications that often occur without prior user expectation.

User actions and eye-tracking data throughout the tasks were logged to a remote server in real-time. This logging included timestamps of each PI display, gaze fixation, duration, and the user's interactions with the task at hand. Such data are collected and subsequently analyzed to understand whether the PIs were noticed and how the interaction context and cognitive load influence their visibility and the user's attention engagement.




Figure 5: Two of the experiment tasks requiring low spatio-temporal attention (*Quiz*) on the left and high spatio-temporal attention (*Game*) on the right.

3.3 Pilot Study

Before running our experiment, we conducted a pilot study with 8 participants in an offline setting. Such a study was promoted from March 22, 2023, to April 13, 2023. We tested the previously described experimental protocol, and we checked whether the tasks were feasible, whether the instructions were sufficient and explicit, whether the web app worked on different devices and their usability, and whether users experienced problems or had doubts during execution. None of the users reported problems in understanding and completing the tasks.

3.4 Data Availability

The data supporting the findings of this study, derived from observations collected as participants performed the tasks, are publicly available ¹.

¹<https://github.com/rmilanese99/privacy-indicators-eye-tracking/>

4 TECHNICAL IMPLEMENTATION OF THE USED SYSTEM

In this section, we describe the development of a web application purpose-built for conducting an experimental evaluation of Privacy Indicators (PIs) and two novel alternatives. The application leveraged the Angular framework for front-end user interaction and WebGazer.js for eye-tracking capabilities. The web app served as a testbed to assess the current PI design's effectiveness and gauge user response to newly proposed PI designs under real-world conditions. Angular was chosen for its scalability and robustness, ideal for creating a dynamic web application capable of supporting the complex functionalities required by the experiment. The application was structured to handle various tasks and conditions systematically, guiding participants through the experiment's flow while capturing their interaction data with high fidelity. The web application's design facilitated a responsive and intuitive user experience, essential for maintaining participant engagement throughout the experiment.

4.1 Integration of Eye-Tracking with WebGazer.js

WebGazer.js (Papoutsaki et al., 2016) provides a versatile eye-tracking solution adaptable to different experimental scenarios. The library was integrated into the Angular application to enable real-time gaze tracking using participants' webcams, with specific focus on the accuracy and responsiveness necessary for tracking quick eye movements.

4.1.1 Calibration Process

The eye-tracking calibration process was a multi-step procedure designed to adapt to the user's unique physiology and environmental conditions. Calibration accuracy was enhanced through iterative user engagement with on-screen elements, which refined the predictive model of WebGazer.js in real-time. The web application was designed to provide visual cues that guide users through calibration without overwhelming or distracting from the core tasks. We used an improved version of Webgazer.js that was calibrated continuously throughout the session. This ongoing calibration, known as drift correction, is dynamically adjusted based on user interactions, specifically their touch inputs. Authors in (Papoutsaki et al., 2018) paper provided a foundational basis for understanding how gaze behavior can be effectively tracked and ad-

justed in real-time, enhancing the accuracy and reliability of eye-tracking data, especially in dynamic user interaction scenarios. This continuous calibration method is crucial for maintaining the precision of gaze tracking throughout the session, ensuring that the data collected reflects the actual user behavior and interactions. This was achieved through minimalistic design choices and using Angular's dynamic rendering capabilities to update the interface based on the progress of the calibration process.

4.2 Privacy Indicator Detection Mechanism

The 'alert zone' setup involved defining the perimeter within the application view where user attention was considered focused on the PI. This logic was encapsulated within the `setPrivacyIndicator()` function, which dynamically adjusted the sensitivity based on the screen's real estate and the user's distance from the device. A dedicated Angular service was implemented to handle the real-time streaming of gaze data provided by WebGazer.js. The service utilized an event-driven model to capture and process gaze coordinates at pre-defined intervals, ensuring the system's responsiveness and the accuracy of the collected data. All gaze data associated with PI detection were timestamped and recorded within the user's session without storing any video feed. This approach was designed to respect user privacy and comply with ethical standards for research involving human subjects.

5 ANALYSIS RESULTS

We analyzed the eye-tracking data collected during the experiment to investigate the effectiveness of different PIs in attracting users' focus during tasks requiring different levels of attention. We introduced two independent variables: the animation extent of the PI (*Classic*, *Edge*, and *Disk*) and the level of concentration required for the task (*Quiz* and *Game*). The dependent variable concerns the eye movement's location and speed as a proxy of attention and awareness of the PI. Specifically, we determined the amount of time spent looking at the PI by measuring the total duration of fixations within a 250 pixel radius around it. We also defined unintended or false fixations as a total cumulative fixation duration of less than 150 ms and excluded them from subsequent analyses. These criteria were adopted to account for the noise and variability introduced by our smartphone-based eye-tracking approach. Out of 174 observations, 3 were discarded based on these criteria, while a further 37

were excluded because the participant completed the task too quickly to achieve the required level of attention, and often before the PI was displayed.

To answer our RQ, we conducted a two-way mixed-effects ANOVA, which tests for the presence of a statistically significant effect of the independent variables, singly or in interaction, on the dependent variable. Given a significant effect, we performed post-hoc comparisons to understand which specific type (or interaction of types) of task and/or PI successfully affected PI awareness. Specifically, we used Bonferroni-corrected unpaired t-tests for pairwise comparisons between different conditions. Importantly, the advantage of this mixed model over alternative solutions is that it takes into account the effect of between-participant variability. Furthermore, the interaction between PI and task allows us to investigate whether a PI might have a specific effect only on one particular task condition and not on the other. To sum up, the null hypothesis (H_0) is that there is no association between the type of task and PI and the awareness of the PI. Conversely, the alternative hypothesis (H_1) is that the presence of PI is noticed differently according to its visual properties and the characteristics of the attentional context. We reject the null hypothesis if the p-value is lower than or equal to 0.05.

In Table 1 and Figure 6 we report the average time spent looking at the PI in the different conditions. These data show a strong effect of task type on PI awareness. Almost all participants didn't fixate on the PI when the task required a strong attentional focus on a different location of the screen (i.e., the *Game* condition). This effect appears to be irrespective of PI animation, although the *Disk* was marginally more effective than the others. Since PIs were overlooked in the *Game* condition, the effect of the different animations can be appreciated only in the *Quiz* condition. *Disk* fixations lasted longer than those of the *Classic* PI. Interestingly, the *Edge* PI was almost not noticed in this condition as well. We attribute this result to the smaller size of the stimulus and its location to the outer periphery of the screen, rather than to its animation type.

Table 1: Average time spent looking at the PI in the different conditions.

		Indicator		
		<i>Classic</i>	<i>Edge</i>	<i>Disk</i>
Task	<i>Quiz</i>	965.81 ms	11.08 ms	2114.40 ms
	<i>Game</i>	0.0 ms	0.0 ms	224.50 ms

The statistical findings are presented in Table 2. The ANOVA corroborated our observations, as we identified significant ($p < 0.05$) main effects of task

and PI type on the interaction. This outcome indicates that our manipulations were effective in influencing PI awareness. Given the lack of fixations in the *Game* setting, post-doc comparisons focused on the *Quiz* one. Unsurprisingly, all PIs had a significant effect, with the *Disk* being the most effective, and the *Edge* being the least effective.

Overall, these results indicate that the PIs were noticed almost exclusively in the *Classic* and *Disk* form during the *Quiz* task. To further investigate the attentional dynamics of these conditions, we measured the time elapsed between the onset of the PI and the first fixation within the PI area. These saccadic reaction times were shorter for the *Disk* PI (mean = 1799; std = 381) compared to the *Classic* one (mean = 2244; std = 513), suggesting that our animation provides a mean for a faster attentional capture. However, an unpaired t-test revealed no significant difference between the reaction times to the two PIs.




Figure 6: Average fixation durations on Privacy Indicators (*Classic*, *Edge*, *Disk*) during low (*Quiz*) and high (*Game*) spatio-temporal attention tasks, highlighting the differential impact of task type on user's attention to PIs.

Table 2: Statistical results.

Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
Type	30163156	2	15081578	9.939948	9.71E-05
Task	29526910	1	29526910	19.46056	2.16E-05
Type:Task	19506429	2	9753215	6.428137	0.002187
Error	1.79E-07	0	0		
Total	2.84E+08	133			

5.1 Discussion

In this section, we discuss the eye-tracking experiment's outcomes regarding the effectiveness of PIs in different task contexts.

The data collected indicate that the cognitive load required by a specific task notably impacts the likelihood of the user detecting a PI. This was particularly

evident in the *Game* setting, which is characterized by a high degree of spatio-temporal focus, and where both Android's native PI, as well as alternative animations, were consistently overlooked. These findings suggest inattentive blindness, where users caught up in demanding tasks may fail to notice secondary stimuli, such as the PIs in this study. Consequently, a key design goal is to develop PIs that can successfully capture the user's attention without interfering with their activity.

There is also evidence that the characteristics of each PI have a direct impact on its visibility, as our analysis revealed that the *Disk* was marginally more effective at capturing user attention than the other PIs examined in this study. This suggests that specific visual features, such as variations in color or motion, may be critical in increasing their detectability, and confirms the need for extensive research to determine the optimal design parameters in different usage contexts.

The implications of our findings regarding the effectiveness of PIs have a significant impact on the privacy of Android users. The *Classic* PI, as implemented, fails to capture the user's attention consistently during tasks requiring high levels of concentration, so malicious apps could exploit this flaw to conceal privacy-damaging actions during periods of intense user engagement. Therefore, there is an urgent need to design PIs that are more noticeable and can intelligently adapt to the current level of cognitive load, ensuring that users are alerted to privacy risks even when completely absorbed in their activity. This aims to safeguard user privacy using instinctive and non-disruptive methods, avoiding the danger of desensitization that can occur with frequent alerts.

While the study provides novel insights into PI effectiveness, it also presents limitations, including the controlled environment of the experiment and the specific demographic of the participants, which may not fully reflect the diversity of real-world application users. The experimental evaluation primarily involved university students, who are generally young and well-educated. This demographic, while reasonable for initial testing, does not encompass all possible user groups, particularly those of different ages, educational backgrounds, and digital literacy levels. Recognizing this, the need to conduct field studies becomes more evident, as such studies could provide a more comprehensive understanding of the effectiveness of privacy indicators across a broader and more diverse population. This would help to reinforce these findings in everyday contexts and ensure the applicability of our proposals to a wider range of users.

6 CONCLUSIONS AND FUTURE DIRECTIONS

Our study contributes significantly to understanding the effectiveness of Privacy Indicators (PIs) in mobile applications. Our findings suggest that the current implementation of the *Classic* PI in Android is inadequate for consistently capturing user attention, particularly in situations that demand high cognitive engagement. This raises concerns about user privacy, as these indicators are the primary alert mechanism against unauthorized resource usage like camera or microphone access by potentially malicious applications.

Exploring alternative PIs, such as *Disk* and *Edge*, revealed that animation and visibility play crucial roles in attracting user attention. The *Disk* PI, with its dynamic animation, demonstrated a higher efficacy in catching the user's gaze in low-attention scenarios compared to the *Classic* and *Edge* PIs. However, the effectiveness of these indicators in high-attention tasks remains limited, indicating a need for more innovative approaches.

Future research should focus on developing more effective PIs that can alert users even during high-concentration tasks. Exploring dynamic and context-sensitive PIs that adapt their visibility based on the user's current activity may prove beneficial. Additionally, investigating user's cognitive load and attention span in relation to PI visibility could offer deeper insights into designing more effective privacy indicators. Further studies could also explore the integration of auditory or haptic feedback as supplementary or alternative alert mechanisms.

An intriguing direction for these studies would be to examine how privacy indicators directly linked to specific in-app actions, such as submitting responses in a quiz or completing a level in a game, might influence users' awareness and attitudes towards privacy. Such investigations could deepen our understanding of the interplay between privacy notifications and user experience in mobile applications, enriching the ongoing dialogue on digital privacy. The ultimate goal is to strike a balance between ensuring user privacy and maintaining a non-intrusive user experience.

ACKNOWLEDGEMENTS

This work has been funded by the European Union - NextGenerationEU under the Italian Ministry of University and Research (MUR) National Innovation Ecosystem grant ECS00000041 -VITALITY-CUP E13C22001060006.

This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 118/2023.

REFERENCES

- Anderson, B. B., Jenkins, J. L., Vance, A., Kirwan, C. B., and Eargle, D. (2016). Your memory is working against you: How eye tracking and memory explain habituation to security warnings. *Decision Support Systems*, 92:3–13.
- Arianezhad, M., Camp, L. J., Kelley, T., and Stebila, D. (2013). Comparative eye tracking of experts and novices in web single sign-on. In *Proceedings of the Third ACM Conference on Data and Application Security and Privacy*, CODASPY '13, page 105–116, New York, NY, USA. Association for Computing Machinery.
- Carter, B. T. and Luke, S. G. (2020). Best practices in eye tracking research. *International Journal of Psychophysiology*, 155:49–62.
- Egelman, S. (2013). My profile is my password, verify me! the privacy/convenience tradeoff of facebook connect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, page 2369–2378, New York, NY, USA. Association for Computing Machinery.
- Elbitar, Y., Schilling, M., Nguyen, T. T., Backes, M., and Bugiel, S. (2021). Explanation beats context: The effect of timing & rationales on users' runtime permission decisions. *USENIX Security'21*.
- Fasano, F., Guerra, M., Milanese, R., and Oliveto, R. (2023). A dynamic approach to defuse logic bombs in android applications. In *Data and Applications Security and Privacy XXXVII: 37th Annual IFIP WG 11.3 Conference, DBSec 2023, Sophia-Antipolis, France, July 19–21, 2023, Proceedings*, page 358–365, Berlin, Heidelberg. Springer-Verlag.
- Felt, A. P., Ha, E., Egelman, S., Haney, A., Chin, E., and Wagner, D. (2012). Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–14.
- Furman, S. and Theofanos, M. (2014). Preserving privacy – more than reading a message. In Stephanidis, C. and Antona, M., editors, *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*, pages 14–25, Cham. Springer International Publishing.
- Guerra, M., Milanese, R., Oliveto, R., and Fasano, F. (2023). Rpdroid: Runtime identification of permission usage contexts in android applications. In *Proceedings of the 9th International Conference on Information Systems Security and Privacy - ICISSP*, pages 714–721. INSTICC, SciTePress.
- Guerra, M., Scalabrino, S., Fasano, F., and Oliveto, R. (2023). An empirical study on the effectiveness of privacy indicators. *IEEE Transactions on Software Engineering*, 49(10):4610–4623.
- Miyamoto, D., Iimura, T., Blanc, G., Tazaki, H., and Kadobayashi, Y. (2014). Eyebit: Eye-tracking approach for enforcing phishing prevention habits. In *2014 Third International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, pages 56–65.
- Papoutsaki, A., Gokaslan, A., Tompkin, J., He, Y., and Huang, J. (2018). The eye of the typer: A benchmark and analysis of gaze behavior during typing. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, ETRA '18, New York, NY, USA. Association for Computing Machinery.
- Papoutsaki, A., Sangkloy, P., Laskey, J., Daskalova, N., Huang, J., and Hays, J. (2016). Webgazer: Scalable webcam eye tracking using user interactions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3839–3845. AAAI.
- Peruma, A., Palmerino, J., and Krutz, D. E. (2018). Investigating user perception and comprehension of android permission models. In *Proceedings of the 5th International Conference on Mobile Software Engineering and Systems*, pages 56–66.
- Punde, P. A., Jadhav, M. E., and Manza, R. R. (2017). A study of eye tracking technology and its applications. In *2017 1st International Conference on Intelligent Systems and Information Management (ICISIM)*, pages 86–90. IEEE.
- Scoccia, G. L., Malavolta, I., Autili, M., Di Salle, A., and Inverardi, P. (2021). Enhancing trustability of android applications via user-centric flexible permissions. *IEEE Transactions on Software Engineering*, 47(10):2032–2051.
- Shen, B., Wei, L., Xiang, C., Wu, Y., Shen, M., Zhou, Y., and Jin, X. (2021). Can systems explain permissions better? understanding users' misperceptions under smartphone runtime permission model. In *USENIX Security Symposium*, pages 751–768.
- Wang, Y., Wang, Y., Wang, S., Liu, Y., Xu, C., Cheung, S.-C., Yu, H., and Zhu, Z. (2021). Runtime permission issues in android apps: Taxonomy, practices, and ways forward. *arXiv preprint arXiv:2106.13012*.
- Whalen, T. and Inkpen, K. (2005). Gathering evidence: use of visual security cues in web browsers. In *Proceedings of Graphics Interface 2005*, GI 2005, pages 137–144, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. Canadian Human-Computer Communications Society.
- Wijesekera, P., Baokar, A., Tsai, L., Reardon, J., Egelman, S., Wagner, D., and Beznosov, K. (2018). Dynamically regulating mobile application permissions. *IEEE Security & Privacy*, 16(1):64–71.