Don't Interrupt Me – A Large-Scale Study of On-Device Permission Prompt Quieting in Chrome

Marian Harbach, Igor Bilogrevic, Enrico Bacis, Serena Chen, Ravjit Uppal,
Andy Paicu, Elias Klim, Meggyn Watkins, Balazs Engedy

Google
{mharbach, ibilogrevic, enricobacis, sereeena, andypaicu, elklm, meggynwatkins, engedy}@google.com

***Abstract*—A recent large-scale experiment conducted by Chrome [4] has demonstrated that a "quieter" web permission prompt can reduce unwanted interruptions while only marginally affecting grant rates. However, the experiment and the partial roll-out were missing two important elements: (1) an effective and context-aware activation mechanism for such a quieter prompt, and (2) an analysis of user attitudes and sentiment towards such an intervention. In this paper, we address these two limitations by means of a novel ML-based activation mechanism – and its real-world on-device deployment in Chrome – and a large-scale user study with 13.1k participants from 156 countries. First, the telemetry-based results, computed on more than 20 million samples from Chrome users in-the-wild, indicate that the novel on-device ML-based approach is both extremely precise (>99% post-hoc precision) and has very high coverage (96% recall for notifications permission). Second, our large-scale, in-context user study shows that quieting is often perceived as helpful and does not cause high levels of unease for most respondents.**

I. Introduction

The web today is a powerful platform that enables developers to build immersive and interactive applications, ranging from real-time communications and content creation to cloud-enabled productivity apps. To support such use cases, browsers have evolved to offer web APIs that are able to provide access to powerful and system-level capabilities, such as notifications, precise location, and reading from the clipboard [29]). To limit the potential abuses of such APIs, browsers often require users to explicitly grant websites permission to use them via runtime permission prompts. While such a prompt-based approach can be effective in preventing unintended access to the underlying capability, it does so at an additional cost for users: their browsing experience can get interrupted with *unwanted* prompts at *inconvenient* moments. For the purposes of this work, unwanted interruptions are caused by prompts that do not lead to a permission being granted. Media outlets, browser vendors and users alike often flag this as an important issue [2], [8], [16], [18], [22].

Recently, Bilogrevic et al. [4] addressed the issue of interruptions due to unwanted notification permission prompts. Their solution relied on a "quieter" prompt user interface (UI) for cases in which users are very unlikely to allow sites to send them push notifications. The experimental results, based on an A/B test with 40 million Chrome users, showed that such quieter prompt UIs led to a reduction of up to 30% in unnecessary user actions (denying or dismissing the prompt) while lowering grant rates by less than 5%. A similar approach has subsequently been adopted by Microsoft's Edge browser and is enabled by default for all of its users since 2020 [6].

On Chrome, such quiet prompts were originally shown if any of the two following conditions were met: (1) if the site's average grant rate was in the lowest 5% of all sites' grant rates (i.e., within the 5th percentile), based on Chrome telemetry data; or (2) if in the past the user denied three consecutive notification permission prompts (on any site) within a 28 day window.

However, these conditions have important limitations. First, condition (1) only affects 1-3% of the overall notification permission requests, and condition (2) is only applicable to 14% of users [4], but then disregards any future contextual signals for those users. Hence, although the quiet prompt itself was effective, its reach was limited. Chrome did not have an effective way of deciding when to show the quiet prompt for most users, and was thus unable to reduce interruptions on a large fraction of unwanted prompts. To address this and thus benefit a significantly larger proportion of users, Chrome needs a more comprehensive activation mechanism.

Furthermore, while the previous experiment was able to reliably determine the impact of the quiet prompt on permission action rates (grants, denies, dismissals and ignores), it lacked any form of feedback from users themselves. It remained unclear how users perceived such interventions and prompts, and whether they understood their options to effectively override the quieting when necessary.

This paper addresses these issues by making the following main contributions:

- We describe how the quiet prompt UI in Chrome changed to become more consistent and be even quieter for certain types of notification permission requests.
- We introduce a novel ML-based activation mechanism for prompt quieting, which uses both contextual real-time signals as well as past actions on permission prompts. We train these ML models server-side and deploy them on-device for local inference on Chrome clients.
- We provide a telemetry-based assessment of this new ac-

Network and Distributed System Security (NDSS) Symposium 2024
26 February - 1 March 2024, San Diego, CA, USA
ISBN 1-891562-93-2
https://dx.doi.org/10.14722/ndss.2024.24108
www.ndss-symposium.org

tivation mechanism's efficacy, based on a sample of more than 20 million prompt requests from Chrome on desktop platforms (Windows, macOS, ChromeOS, Linux). We find that Chrome can now mediate 43% of notification and 24% of geolocation permission prompts, thus increasing its impact by more than 10x. The ML-based activation mechanism achieves 99% post-hoc precision and 96% post-hoc recall for the notifications permission. This substantial improvement over the status quo will lead to less unnecessary prompts and can thus reduce prompt blindness and avoiding unintentional grants.
- We conduct the first, to the best of our knowledge, large-scale and in-context user survey with 13.1k participants from 156 countries on attitudes towards prompt quieting for Chrome on desktop platforms. According to the results, 84% of respondents rated quieting as at least moderately helpful, whereas only 10% felt very or extremely uneasy about it. We also find that the quiet prompt UI can be easily ignored and meets the design goal of neither being too noticeable nor not noticeable enough, as 51% of respondents indicate they did not notice the quieted prompt and thus were not interrupted by it. Our findings demonstrate that our intervention on permission prompting resonates with users, as long as a sense of control is maintained.

The remainder of the paper is structured as follows. We provide additional background and related work in Section II, followed by a description of the improved UI design in Section III. Section IV introduces the improved activation mechanism and presents its evaluation based on Chrome telemetry. We describe the in-product survey methodology in Section V, and present the results in Section VI. We summarize the limitations of this work in Section VII and discuss open challenges and how they can be addressed in Section VIII. Finally, we conclude the paper and outline potential next steps in Section IX.

II. Background and Related Work

Our work builds on the prior work of Bilogrevic et al. [4], which laid the foundations for less intrusive web permission prompts in Chrome. In this section, we summarize recent developments regarding notification permissions on the web as well as previous work on evaluating user sentiment of security- and privacy-related UIs.

A. Notification Permission Prompts

Modern software platforms – such as Android, iOS, Windows, macOS, or the web – offer APIs that enable content providers to send push notification messages to users who granted the related permission. Such APIs are designed to be used for sending timely and relevant messages, as they usually interrupt the users' current activities and take up a portion of the display. Due to their significant potential to redirect users' attention towards these notification messages, push notifications are prone to abuse by unscrupulous content providers, who might try to boost the traffic to their properties by sending a large number of irrelevant – or outright abusive – notifications [3], [27].

In an attempt to limit such abuses, Mozilla made changes to how websites can ask for notifications in 2019 [21]. Similarly, in 2020, Chrome created, experimented with and partially released a mechanism to reduce unwanted interruptions due to permission prompts on the web [4], which reduced the unnecessary user actions on permission prompts by 30% while only reducing grant rates by 5%. Following up on that, Microsoft decided to release a similar quiet prompt UI for all notification requests to all of their Edge browser's users [6]. We describe Chrome's quiet prompt UI as well as when it gets activated in more detail in Sections III and IV.

B. Understanding Security & Privacy Decision UIs

As we are currently not aware of other adaptive interventions on when websites or apps get to show permission prompts, we lean on the evaluation of other security and privacy decision UIs. Evaluating security and privacy decision UIs is notoriously difficult, as such decisions are usually heavily context-dependent [1], [19], [23], [24]. Previous work in this space often relied on either elaborate hypothetical setups or field studies to collect data. For example, Bravo-Lillo et al. [7] asked crowd-workers to evaluate online games as a decoy and showed warnings and permission prompts during this task. Similarly, Elbitar et al. [12] presented users with goals unrelated to the study's primary purpose when evaluating permission prompt timing and rationale strings. Several studies used instrumented Android phones to collect permission decision making on prompts or settings in the wild [9], [17], [30]. Harbach et al. [15] used experience sampling to understand phone unlocking behaviors in context, combining telemetry with qualitative feedback via short, in-context surveys.

When it comes to user sentiment of security and privacy decisions or interventions in browsers, Felt et al. [14] evaluated several designs for TLS warnings. They used micro-surveys to measure comprehension before measuring adherence using telemetry. As a follow-up, Reeder et al. [25] used experience sampling via browser extensions for Chrome and Firefox to achieve a compromise between using hypothetical scenarios (lacking ecological validity), and relying on telemetry (lacking deeper, more qualitative insights).

Overall, prior work suggests that in-situ collection of user sentiment and attitudes yields the most valuable data due to its ecological validity. We thus embrace this approach for this study.

III. UI Treatment of Permission Prompt Quieting in Chrome

Based on results from A/B experiments, Bilogrevic et al. [4] originally chose UI patterns for quiet permission prompts (cf. Figures 1 and 2). The quiet prompt UI in desktop Chrome has changed in several ways since it launched in Chrome version M80 [20]. These changes include migrating the quiet prompts to use a new UI pattern more consistent with other permission surfaces and adding a UI treatment that is even less interrupting.

Hereafter, we discuss these changes and the new status quo of quiet prompt UI in Chrome. Note that we focus our discussion in this and the following sections on Chrome for desktop platforms, as we conducted the in-product survey only

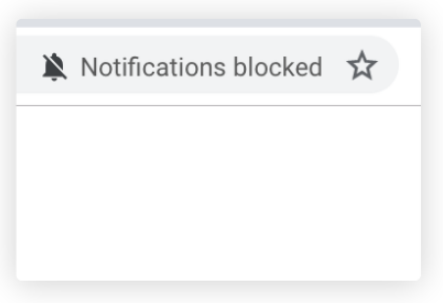

Fig. 1: Quiet permission prompt on Chrome desktop from Bilogrevic et al. [4].

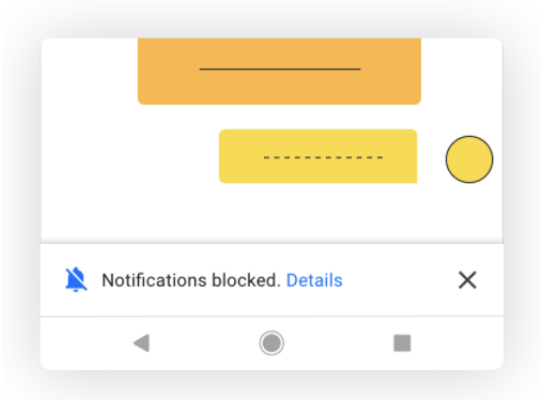

Fig. 2: Quiet permission prompt on Chrome on Android from Bilogrevic et al. [4].

on desktop platforms. Evaluating the status quo of permission prompt quieting on mobile platforms is subject of future work.

A. *Quiet UI consistency*

The original placement of the quiet UI was aligned with other page-related actions (see Figure 1), such as automatic pop-up blocking or turning off access to already granted capabilities, such as camera or location. However, this location to the right of the address bar was also disadvantageous in two important ways. Firstly, quieted prompts appeared in a different location than full prompts, which are aligned to the left-hand side of the address bar. This difference between "normal" and quiet UI likely contributed to the drop in grant rates noted by Bilogrevic et al. [4]. Secondly, the site controls surface is accessed from the left-hand side of the address bar (from what used to be the lock icon and now is the tune icon, [11], [28]). The site controls surface has permanent controls for permissions as well as other security- and privacy-related information. Showing the quiet prompt on the right-hand side was thus a missed opportunity for reinforcing an entry point to permanent permission controls.

In January 2022 with Chrome M97, the quiet prompt started using a new UI pattern on the left-hand side of the address bar, improving locality and consistency with other permission-related UI elements. To avoid cluttering this space with different icons and styles, the Chrome UX team aligned on a chip pattern that can be reused for multiple purposes. Figure 3 shows the current design. The colored background of the chip evokes the style of a button and thus provides a more perceptible interaction affordance than the previous, flat visual (cf. Figure 1). The progressive collapsing behavior gives users an additional chance to notice the movement and grant the permission, if necessary. We refer to this new design and placement as the "*quiet chip*" in the rest of the paper.

B. *An even quieter prompt UI*

Originally, Bilogrevic et al. [4] experimented with both site-based and user-based activation to be shown with the same UI treatment. Leading up to the full launch of this feature with M83 in June 2020, the Chrome team decided to introduce an even quieter treatment, which displayed the collapsed state of the regular quiet prompt immediately. This was first implemented using the old UI pattern on the right-hand side of the address bar and was then also moved to the new, chip-based pattern on the left-hand side (cf. Figure 3). We refer to this as the *quietest* prompt UI or "*quietest chip*". This treatment applies to the site-based activation mechanism described in Section IV, and was introduced due to Chrome's high confidence that users are very unlikely to allow access and thus should experience minimal interruption.

Additionally, in December 2022, the Chrome team identified sites using notifications in disruptive ways (extremely high ratio of notification messages shown and not interacted with per user engagement with the site) using telemetry of opted-in clients. On the top 30 such sites on desktop and mobile, respectively, notification permission requests have since been surfaced using the quietest prompt, warning users of this potentially disruptive behavior. In addition, preexisting notification permissions on these sites, which were granted through a prompt without users being adequately warned, were revoked.

IV. Improved Activation of the Quiet Prompt UI

The prior work of Bilogrevic et al. [4] described and evaluated permission prompt quieting as it initially launched in Chrome version M80 [20]. At that time, quieted prompts were displayed based on the following three activation mechanisms:

1) *Site-based:* Two lists of "interrupting websites", one for desktop and one for mobile clients. These lists are generated based on aggregated Chrome telemetry (at URL origin level) of prompting and granting behaviors. Included sites fall into the bottom 5% of sites in terms of notification permission grant rate.
2) *Opt-in:* User opt-in via a "Use quieter messaging" option in Chrome settings, which quiets all notification prompts.
3) *User-based:* Permanent activation of the quiet treatment on all websites after three consecutive deny decisions in a row on any website within 28 days.

Their evaluation results, based on these mechanisms and the original UI described at the beginning of Section III, were encouraging. The average deny rate across websites in their experiment group decreased by 22.5% on desktop and by 30.0% on Android, suggesting that users needed to make fewer unnecessary decisions. At the same time, the average grant rate

Fig. 3: New quiet chip UI on the left-hand side of the address bar, co-located with other permission surfaces. The quiet prompt is shown on the left above and the additional "quietest" treatment on the right of this figure.

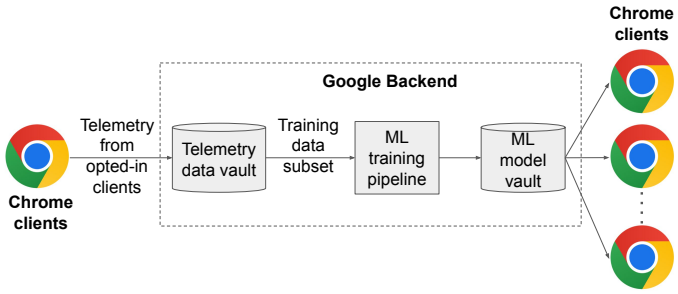

Fig. 4: Web Permissions Prediction Service (WPP) training and deployment architecture.

only decreased very little (by 3.7% and 5.0%, respectively), indicating that users who want to grant can still do so.

However, these activation mechanisms have important limitations. First, the site-based mechanism only affected 1-3% of the overall notification permission requests, and the user-based mechanism was only applicable to 14% of users [4]. Hence, although the quiet prompt itself was effective when it showed, the number of situations where it could be showing was limited.

In the following subsections, we present an improvement over this prior work and discuss its efficacy based on Chrome telemetry data. Sections V and VI then present a user-centered evaluation of the new status quo of prompt quieting in Chrome.

A. ML-Based activation

To address the limitations outlined above, we designed and deployed a novel ML-based mechanism that decides when to show the quiet prompt UI, replacing the previous user-based activation mechanism. Figure 4 shows the high-level architecture of the *Web Permissions Predictions (WPP)* model, which was briefly mentioned in [5] and [13], but never fully described in prior work.

WPP relies on a ML model that is trained on telemetry data that Chrome normally collects from a subset of opted-in Chrome users. In particular, the users contributing training data need to have two settings turned on in Chrome: (1) share usage reports and crash analytics with Google, and (2) "Make searches and browsing better / Sends URLs of pages you visit to Google" [10]. The model is trained on both contextual as well as statistical features from users' past actions on web permission prompts.

Specifically, the features used for training and inference are:

- The permission type, which could be "notification" (i.e., the website would like to send push notification messages to the user) or "geolocation" (i.e., the website would like to know the user's geographical location).
- Average action rates (i.e., grant, deny, dismiss and ignore) across all permissions over the last 28 days on "loud" (i.e., non-quiet) prompts, rounded to the first decimal.
- Average per-permission action rates (i.e., grant, deny, dismiss and ignore) over the last 28 days on loud prompts, rounded to the first decimal.
- The total number of loud permission prompts shown over the last 28 days, bucketized in a non-linear way, capped at 20.
- Whether there was a user gesture (a click anywhere on the content area or keyboard event within 5 seconds) before the permission request.
- The platform, which could be either desktop (e.g., Windows, macOS, ChromeOS, Linux) or mobile (e.g., Android).

Note that the WPP model does not use or process any client identifiers; WPP only uses a subset of the data that users contribute as part of the two Chrome settings described earlier [10]. Furthermore, WPP's scope was expanded to cover the Geolocation API[1] permission as well, which is the second-most frequently asked permission type, according to Chrome telemetry. For privacy and performance reasons, each of the statistical features described above is pre-processed (coarsened or bucketized) before leaving the client. We conducted a series of experiments to assess the impact of coarsening and bucketing on WPP features. Based on these experiments, we decided to round the average action rates to the first decimal place and to bucketize the total number of shown prompts in a non-linear way so that each bucket has a similar number of requests (buckets: $0, 1, 2, 3, 4, 5, 6, 7, 8, 9, [10, 11], [12, 14], [15, 19], [20, +\infty)$). Our results show that this has a minimal effect on prediction accuracy, with more than 99% of the predictions resulting in the same UI being chosen. However, coarsening and bucketing have a significant positive impact on privacy, as the probability of a set of features being unique to a user decreases from 2.5% to 0.2%.

The WPP ML production pipeline relies on the TFX framework[2] to manage data ingestion, pre-processing, training, evaluation, validation and deployment. In particular, the ML model uses a neural network architecture, and is optimized to achieve a precision greater or equal to 95%. This specific value was chosen to ensure a maximum error rate of less than 5% if Chrome were to show a loud UI, and in that case show the quiet UI. As the quiet UI makes the prompt less interrupting and thus is less visible to users, we wanted to limit its error rate in order to reduce the potential negative effect on grant rates

[1] https://developer.mozilla.org/en-US/docs/Web/API/Geolocation_API
[2] https://www.tensorflow.org/tfx

on any given site to at most 5%, when compared to always showing the loud UI.

After the WPP model has been trained, Chrome uses the Component Updater component[3] to deploy the TFLite[4] version of the model to devices.

Currently, WPP is enabled for users who utilize Safe Browsing (Standard or Enhanced). If that is the case, Chrome adopts the following decision logic to select which prompt UI to show (i.e., loud, quiet or quietest) every time a website wants to request either a notifications or a geolocation permission:

1) *Site-based:* If the site is among the sites with the lowest grant rates for that permission, show the *quietest* prompt UI.
2) *Opt-in:* Otherwise, if the user has enabled "Use quieter messaging" in Chrome Settings, show the quiet prompt UI.
3) *ML-based:* Otherwise, if Chrome has shown 4 or more loud permission prompts over the last 28 days and if WPP predicts that the user is very unlikely to grant the permission request, show the quiet prompt UI.
4) Otherwise, show the loud prompt UI.

Note that based on the contribution described above, the user-based activation mechanism has been replaced by the ML-based activation mechanism. Also note that WPP is not automatically triggered for every permission request where it could theoretically apply. That is, Chrome needs to have shown 4 or more loud permission prompts to the user over the last 28 days. The reason for this additional check is to not quiet too many prompts and only start doing so for users that see more than 1 prompt per week, on average. As the number of prompts seen by a user in a given week fluctuates, we use a longer time window.

B. Telemetry Results

Table I shows the performance metrics for the improved activation mechanism based on WPP on desktop platforms. Our evaluation is based on more than 20 million permission prompt requests in June 2023, coming from a sample of Chrome users who opted-in to sending telemetry to Google.

Out of all the samples, we can see that ML-based activation of prompt quieting was considered for 43% of the notification permission prompts and for 24% of the geolocation prompts. The vast majority of those prompts were quieted (96% of the time for the notifications permission and 81% of the time for the geolocation permission). For the remainder of prompts, WPP indicated that the user may be likely to grant these permissions and the prompt was thus not quieted.

The post-hoc precision value (99%), which is computed on actions users took on the quieted prompts after they were shown, indicates that WPP had a less than 1% false positive rate. This means that WPP showed a quiet prompt and the user subsequently granted this request in 1% of cases. Note that post-hoc precision measures both the *intent* and *ability* to grant using the quiet prompt UI, and thus a lower post-hoc false positive rate (1%) is expected when compared to the pre-hoc false positive rate during the training process (5%), due to the smaller visual footprint of the quiet UI (post-hoc) as compared to the loud UI (pre-hoc) that was used for training. Similarly, the high post-hoc recall value, which is computed on all prompts for which WPP was the UI selector, shows that WPP was able to correctly show the quiet UI for the vast majority of prompts that were not granted, thus reducing the unwanted requests while still showing the loud UI for the ones that users wanted to grant.

TABLE I: Performance metrics for WPP on desktop platforms, as reported from telemetry.

Metric	Notifications Permission	Geolocation Permission
# of prompts	$>$ 10 million	$>$ 10 million
% of prompts for which WPP was the UI selector	43%	24%
% of quieted prompts (over all prompts for which WPP was the UI selector)	96%	81%
Post-hoc precision	99%	99%
Post-hoc recall	96%	83%

V. In-product Survey Methodology

The results of Bilogrevic et al. [4] as well as what we presented in Section IV-B above show encouraging user behavior on quieted prompts. To ensure we are serving users' needs well, we wanted to additionally understand user attitudes towards such an intervention as a whole. We thus designed and conducted a user study, collecting user feedback whenever Chrome shows a quiet or quietest chip. This comprises all mechanisms for triggering quieted prompts shipped in Chrome M111, i.e. the site-based, opt-in, and ML-based mechanisms described in Section IV. The study aimed to answer the following research questions:

RQ1. Are quieted prompts easy to ignore and are neither too noticeable nor not noticeable enough?
RQ2. To what extent do users have useful intuitions about permission interventions?
RQ3. Does quieting provide user value?
RQ4. To what extent are users concerned about the intervention and, if so, what causes that concern?
RQ5. To what extent do users agree with the quieting decisions?
RQ6. Do users know how to override quieting?
RQ7. Do users know how to disable quieting?

The study was conducted as follows. We launched two in-product surveys in Chrome, to keep each individual survey short. Both surveys addressed RQ1, as we wanted to begin with a simple question and allow at least a rough comparison of the samples obtained between surveys. Then, only the first survey covered RQs 2 – 4, whereas only the second survey covered RQs 5 – 7. The full questionnaire can be found in Appendix A. Note that two questions (Q3 and Q5) in survey 1

[3]https://chromium.googlesource.com/chromium/src/+/lkgr/components/component_updater/

[4]https://www.tensorflow.org/lite

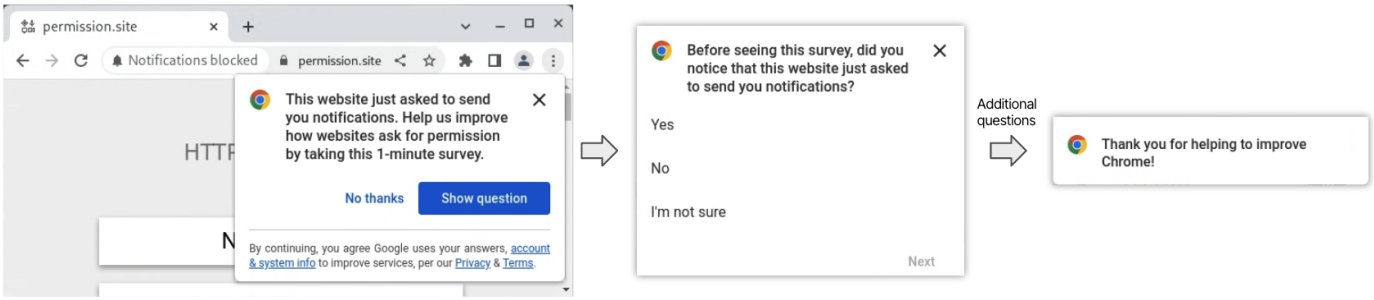

Fig. 5: Screenshot of survey invitation and subsequent screens. The final 'Thank you' message disappeared automatically after 5 seconds.

TABLE II: Response behaviors across the six surveys we showed. s1 and s2 denote the two individual surveys per intervention. "# partial" denotes incomplete responses. Some surveys overshot their target response counts. We discarded the additional responses and final response counts can be found in the "retained" column.

Intervention – Survey	# prompts	% accepted	# partial	% partial	# completed	# retained	median time	# EN	# ES	# FR	# JA	# br-PT	# RU	# zh-CN
geolocation s1	530,349	1.3%	3,661	53.9%	3,127	2,182	63	1,233	239	50	62	347	217	34
geolocation s2	388,444	1.3%	2,574	51.7%	2,407	2,197	66	1,221	278	67	70	371	166	24
notifications s1	880,593	1.4%	6,020	49.9%	6,033	2,174	69	1,001	256	98	234	372	174	39
notifications s2	287,470	1.4%	1,833	45.5%	2,195	2,187	68	960	301	82	285	386	135	38
quietest notif. s1	467,980	1.1%	1,867	36.8%	3,209	2,182	94	1,438	120	42	39	38	488	17
quietest notif. s2	313,463	1.1%	1,377	38.6%	2,192	2,187	97	1,305	94	38	27	48	657	18
	2,868,299	1.3%	17,332	47.5%	19,163	13,109	66	7,158	1,288	377	717	1,562	1,837	170

gave a brief explanation of why the website's request was blocked. This explanation was adjusted to match the reason for showing the respective UI treatment and thus let us understand if users have different attitudes towards quieting if it is done for different reasons. For the quiet chip, it said *"based on your past choices"*. For the quietest chip, it said *"because most people block it or notifications from this site may be disruptive"*. Furthermore, questions Q3 and Q4 in survey 2 provided answer options for possible override and disabling actions. The answer options offered were selected based on a cognitive walkthrough of the UI, where UX experts determined the most likely actions users may take. We prioritized a concise description of possible actions over making it harder to spot "correct" answers (both of which coincidentally begin with "Click on..."). To counteract anchoring effects further, answer options on these questions were shown in randomized order.

Due to technical constraints, each of the two surveys was fielded separately on three different interventions: Notification requests using the quiet chip, Geolocation requests using the quiet chip, and Notification requests using the quietest chip. Therefore, there was a total of six independent surveys (two surveys by three intervention types). The surveys were in field simultaneously for two weeks in March 2023, after Chrome M111 became available[5]. The survey was only available on Chrome desktop due to another technical limitation, which currently prevents us from timing surveys to be shown after permission prompt interactions on mobile platforms.

Chrome users are eligible to see an in-product survey when all of the following conditions for their Chrome profile are met:

- Not opted out of "Help improve Chrome's features and performance" setting;
- Not displayed another in-product survey within 180 days;
- Created profile or installed Chrome at least 30 days ago;
- Chrome is not recovering from a crash; and
- A survey language matching the current Chrome language (locale) is available.

One of the two surveys would randomly show approximately five seconds after a quiet or quietest prompt starts showing. This delay is due to a technical limitation where surveys cannot be pre-fetched because of server load constraints. If the conditions for showing a survey were met, a user would first see an invitation page, and then one question per page (cf. Figure 5). Respondents were able to abandon the survey at any time by clicking the "x" button in the top-right corner. The quiet or quietest chip UI would keep showing as long as the survey was showing in order to help respondents understand what it is about.

1) Ethical Considerations: We are not subject to IRB review, however a cross-functional team of stakeholders as well as other user experience (UX) researchers at Google reviewed and approved the research plan. All UX researchers received formal training on research ethics and we followed standard company practices for ethical user research. We further did not retain any identifying data with our survey. In particular, while IP addresses were transiently used to determine respondents' likely countries of origin, only the country was retained and associated with the survey response.

Survey participation in Chrome is only offered to users at most once per 180 days and only if they did not opt out of helping to improve Chrome. The survey itself was short and

[5]During response collection, there was a problem for two of the six surveys in Russian language and they did not collect any responses initially. These responses were back-filled in early April 2023.

easy to ignore or dismiss. The survey invitation provided links to our privacy policy as well as an overview of any additional data sent along with their responses. This data comprised which permission type they saw a prompt for, their user agent string, a timestamp and their timezone offset.

2) Translation: The survey was originally written in English, and subsequently translated into six other languages: Spanish, French, Japanese, Brazilian Portuguese, Russian, and Chinese. They were selected to cover the locales that contribute at least 3% of all Chrome profiles reporting telemetry. Table II shows response counts by locale, intervention type and survey. Translations were created by professional translators, who also translate all other texts in Chrome. These translations were then reviewed by at least 1 native speaker of that language at Google who also use English at work on a daily basis. Translation issues were flagged and resolved with the original translators.

3) Responses: We collected 2,200 completed responses for each (intervention type, survey) pair, to ensure sufficient participation across all locales. Table II shows an overview of response behaviors. The accept and abandonment rates are consistent with other such surveys we field in Chrome.

It is noteworthy that we see less abandonment for surveys showing after users encountered the quietest chip treatment of notification prompts. Respondents in this condition also took substantially longer to answer, possibly because it might have been more difficult to understand what was happening, based on the smaller UI surface area of the quietest chip. Finally, with the quietest chip on notifications, the fraction of respondents answering in Russian language is also substantially increased. This suggests that the intervention may be more frequently triggered for respondents using a Russian locale.

TABLE III: Top 10 respondent countries based on their IP address

Country	# Responses
USA	3,509
Brazil	1,586
Russia	1,146
India	793
Japan	732
Canada	499
Mexico	359
Ukraine	349
Spain	296
France	246

In total, across all six surveys, we had representation from 156 countries (at least one respondent from each country, per geo-coded IP addresses). Table III lists the top ten countries by respondent count. Unfortunately, due to the nature of the in-product surveys, we do not have any additional details about respondent demographics.

4) Open-ended Responses: Participants that reported at least moderate levels of unease with Chrome's intervention in response to the question asked about RQ4 (see Section VI-C) were invited to elaborate on this sentiment in an open-ended question. Responses languages other than English were machine-translated to English using Google Translate and first inductively coded by one of the authors. A second author then independently coded all responses using the same code book again (initial Cohen's $\kappa = .53$). Conflicts were resolved in discussion. Respondents that provided unintelligible or nonsensical open-ended answers were removed from the data set. The "retained" column in Table II shows the final count of respondents.

5) Statistical Testing: To compare response proportions between various slices of the data, we use omnibus χ^2 tests and report pairwise differences when the absolute value of standardized residuals ($sresid$) is at least two [26]. While we are not testing hypotheses involving differences based on survey locale, we note statistically significant differences between survey locales for each analysis we report, to make differences potentially caused by diverging response behaviors transparent.

VI. In-product Survey Findings

This section will complement the telemetry results presented in Section IV-B with users' attitudes and perceptions. We walk through results for each research question in the subsections below.

A. Ease of Ignoring and Noticing the Prompt (RQ1)

For both surveys, the first question after accepting the survey invitation captured if respondents noticed the quiet prompt UI in the address bar. We wanted to avoid asking hypothetical questions to understand if prompts are easy to ignore. We assume that if respondents didn't see the quieted prompt before the survey, they would have ignored it during their regular browsing as well. Beyond that, even if users noticed the prompt, they may still choose to ignore it. Note that, while the survey was active, the chip would keep showing to give users a chance to understand what the survey is asking about, further increasing the chance of respondents noticing it.

Overall, 51% of respondents report not noticing the quiet UI variants. Given that we neither want the prompt UI to be too noticeable nor not noticeable enough to avoid undue interruptions, we consider this finding encouraging. From a behavioral point of view, telemetry results (cf. IV-B) show that the quiet prompt UI gets ignored 99% of times it was shown. In alignment with the design goal of not interrupting users, this confirms that most users do not act on the quiet prompt UI. Nevertheless, users should be able to override if they need to, which will be looked at in the following subsections.

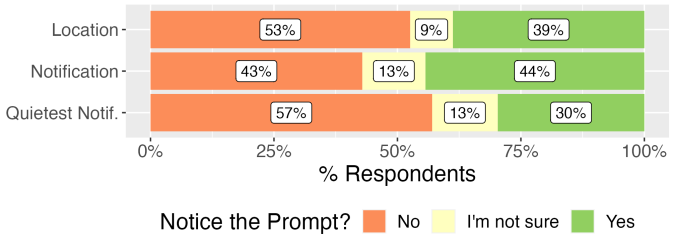

Fig. 6: Respondents' answers on noticing the quieted permission prompt split by intervention type.

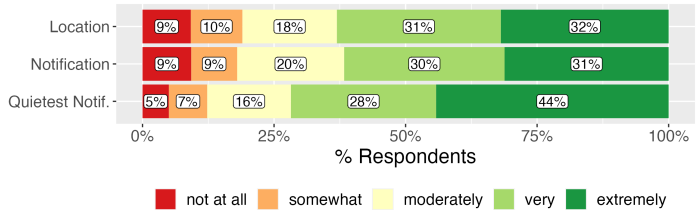

(a) **Helpful** – Q: *Chrome automatically blocked this website's request, based on your past choices. How helpful do you find Chrome's action?*

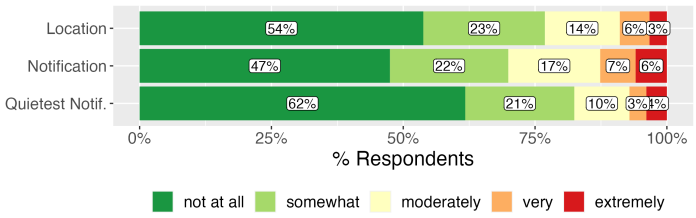

(b) **Uneasy** – Q: *How uneasy do you feel about Chrome's action?*

Fig. 7: Respondents' ratings of how helpful Chrome's prompt intervention is and to what extent it makes them feel uneasy.

As Figure 6 shows, the quiet prompt for location requests went unnoticed more frequently than notification requests ($\chi^2(4) = 265.4$, $p < .001$, $sresid(No, Notification) = -7.4$). Additionally, the quietest notification version remained unnoticed even more frequently ($sresid(No, QuietestNotification) = 5.7$). We speculate that location is less valuable when using websites on a Desktop computer and thus it is less expected that a website would ask for it. The quietest treatment of notification appears least noticeable, likely because the UI surface area is simply smaller.

There was no difference between response bucket proportions between the two surveys across the intervention types. The proportion of respondents not noticing the prompt was higher with a Russian locale ($\chi^2(6) = 139.5$, $p < .001$, $sresid = 2.6$) and Japanese ($sresid = 4.2$) locale and lower for Spanish ($sresid = -5.6$) and French ($sresid = -2.6$).

B. Understanding the Intervention (RQ2)

Next, respondents taking survey #1 were told that this website's permission request was blocked and asked why they think this happened. It is important to note that there was no prior marketing or user education on this feature, so we cannot expect users to have a clear mental model of what is happening. This question aimed to understand if there is sufficient intuition about the intervention.

TABLE IV: Responses on why this website's request was quieted. Bold text indicates answers with what we deem an "appropriate" intuition.

Reason	Geolocation	Notif.	Quietest Notif.	Total
Chrome thinks that this website is dangerous	15.1%	13.4%	14.9%	14.4%
Chrome thinks that I'm not interested in this website	4.2%	8.0%	9.7%	7.3%
I don't know	50.0%	46.3%	40.1%	45.5%
Previously denied request	16.2%	17.9%	19.0%	17.7%
Told Chrome to block website	9.4%	10.4%	12.6%	10.8%
Other	3.6%	2.1%	2.4%	2.7%
This website has a technical issue	1.6%	2.0%	1.3%	1.6%

As shown in Table IV, respondents' intuition was accurate in 22% of cases. Almost half of respondents plainly indicate that they do not know. For the quietest treatment of notification requests, respondents were significantly more likely to select "Chrome thinks I'm not interested" ($\chi^2(12) = 102.5$, $p < .001$, $sresid = 4.1$) and "I have told Chrome to block this website" ($sresid = 2.5$). In turn, respondents in this condition were also less likely to respond with "I don't know" ($sresid = -3.7$).

Across locales ($\chi^2(36) = 298.0$, $p < .001$), those responding on a Russian locale were more likely to think "I have told Chrome to block" (22.2% vs. 10.8% overall, $sresid = 10.3$), those on a Spanish locale were more likely to indicate "Chrome thinks this website is dangerous" (18.0% vs. 14.4% overall, $sresid = 2.4$), and those on an English locale were more likely to select "I previously denied" (19.9% vs. 17.7% overall, $sresid = 3.3$). "I don't know" was selected more frequently by those answering in Japanese (59.4% vs. 45.5% overall, $sresid = 3.8$) and Portuguese (51.7% vs. 45.5% overall, $sresid = 2.5$).

C. Helpfulness & Unease (RQs 3 & 4)

Respondents of the first survey were asked to rate how helpful they find Chrome's intervention as well as to what extent this makes them feel uneasy. We chose helpfulness to operationalize user value more generally, as other, more specific value propositions (like "not being interrupted") appeared too hard to reason about for respondents. Similarly, we chose ratings of feeling uneasy to operationalize being concerned in a more general way.

84% of respondents found quieting at least moderately helpful, with 66% even finding it very or extremely helpful. As Figure 7 shows, this is more pronounced for the quietest treatment on notifications (88%, $\chi^2(2) = 41.4$, $p < .001$, $sresid = 2.1$). At the same time, 24% felt at least moderately uneasy about the quieting intervention overall, but only 10% rated this as "very" or "extremely". While participants reported more unease for regular quieting of notification requests (30% at least moderately uneasy, $\chi^2(2) = 95.1$, $p < .001$, $sresid = 6.2$), there was substantially less such feeling reported for the quietest treatment (18%, $sresid = -5.8$).

This difference may be explained by differences in how we explained why prompts were quieted in the survey. While the quiet treatment explanation mentioned quieting happening "based on your past choices", the explanation for the quietest version mentioned "because most people block it or notifications from this site may be disruptive". These findings suggest that quieting is perceived as even more helpful and less concerning when using these criteria.

TABLE V: Codes assigned to open-ended responses on reasons for feeling uneasy about Chrome's intervention.

Reason Category	Example	Geolocation	Notification	Quietest Notif.	Total
Want more control	*should ask first, make recommendation instead, feels like censorship*	51 (29%)	47 (19%)	41 (24%)	139 (23%)
Unsure what is happening	*general confusion / want to know more*	20 (11%)	28 (11%)	15 (9%)	63 (11%)
Inappropriate blocking in this case	*doesn't make sense on the this site, can't be perfect*	11 (6%)	33 (13%)	14 (8%)	58 (10%)
Fear of missing out	*afraid to miss something, may change their mind*	10 (6%)	25 (10%)	12 (7%)	47 (8%)
Privacy	*Chrome knows too much*	10 (6%)	18 (7%)	9 (5%)	37 (6%)
Concerned about malware/hackers	*site is not safe*	6 (3%)	16 (7%)	1 (1%)	23 (4%)
Unclear or off topic		29 (17%)	40 (16%)	34 (20%)	103 (17%)
No concern/probably OK		13 (7%)	18 (7%)	26 (15%)	57 (10%)
Answered unease question in reverse		8 (5%)	6 (2%)	4 (2%)	18 (3%)
Total		175	246	173	594

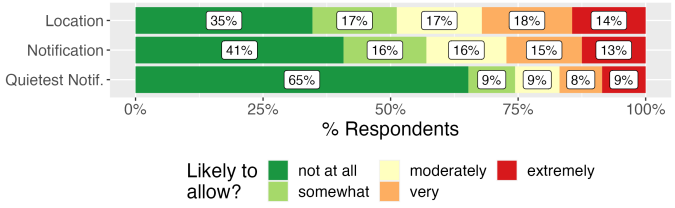

Fig. 8: Respondents' ratings of their likelihood of allowing the current website to use the given capability after all.

One additional aspect to note is that 3% of respondents indicated in their open-ended response to have misunderstood the unease question and its response scale, answering it in reverse (cf. Table V).

Between locales, those answering in Russian were slightly less likely to find quieting at least moderately helpful (77%, $\chi^2(6) = 78.5$, $p < .001$, $sresid = -2.3$). Perceived unease varied more substantially between survey locales: Japanese (56% at least moderately uneasy vs. 24% overall, $\chi^2(6) = 310.1$, $p < .001$, $sresid = 12.0$) and Portuguese (33% at least moderately uneasy, $sresid = 5.3$) locales report more concern while those answering in Russian (13% at least moderately uneasy, $sresid = -6.4$) and Chinese (2% at least moderately uneasy, $sresid = -4.2$) report lesser amounts.

D. Reasons for Feeling Uneasy (RQ4)

Of the $1,543$ respondents reporting at least moderate unease, 594 provided an open-ended response on why they think they feel that way. Table V provides an overview of the reason categories identified during coding (cf. Section V-4) and mentioned at least 15 times across the three intervention types.

A perceived lack of control was the most commonly cited reason. Many respondents also indicated that they weren't sure what was going on and thus stated to feel uneasy about the intervention. Another common reason was perceived inappropriateness of blocking the current site's request. Similarly, several participants also remarked that quieted permission requests could have them miss out on relevant information or functionality. Finally, a few participants also had privacy concerns or were afraid the site itself was malicious and therefore blocked.

E. Subjective Efficacy of Quieting (RQ5)

In the second survey, we asked respondents to rate how likely they are to allow the current website to use the requested capability. In comparison with our telemetry, this allows us to compare objective (behavioral) and subjective (attitudinal) false-positives as a measure of agreement with the quieting decision.

As detailed in Figure 8, 61% of respondents felt less than moderately likely that they would allow. In contrast, 25% indicated being very or extremely likely to allow. Respondents seeing a quietest notification prompt were substantially less likely to want to allow (74% less than moderately likely to allow, $\chi^2(4) = 244.7$, $p < .001$, $sresid = 7.8$). Across survey locales, participants responding to in English reported being less than moderately likely more frequently (68%, $\chi^2(6) = 201.9$, $p < .001$, $sresid = 5.0$), while those responding in Spanish (44%, $sresid = -5.4$) and Japanese (43%, $sresid = -4.1$) did so less frequently.

These findings suggest two interpretations: First, the activation logic for the quietest chip appears to create less subjective false positives, which is congruent with the stricter criteria for block list inclusion. Second, the subjective false-positive findings superficially seem at odds with the findings from our telemetry (99% precision, cf. Section IV-B). However, it is plausible that we are observing another disconnect between intentions and actual behavior (often referred to as the "privacy paradox", [1]).

F. Finding the Escape Hatch (RQ6)

Respondents in the second survey were asked which action they would first try to still allow the website to access to the capability. They were able to choose from a predefined list of four actions, two of which – clicking the lock icon or the chip itself – can be considered "useful", in that they help to make progress towards the goal of overriding the block decision. Additionally, we also offered an "Other" option as well as "I don't know".

As Figure 9 details, only 40% of respondents had a useful intuition about what to do. 32% outright stated that they would not know what to do. Respondents seeing the quietest chip were slightly more likely to take a useful action (43% vs. 38-39%, $\chi^2(2) = 14.4$, $p < .001$, $sresid = 2.3$).

Between survey locales, those responding in Japanese were even more likely to not know what to do (51% vs. 32% overall,

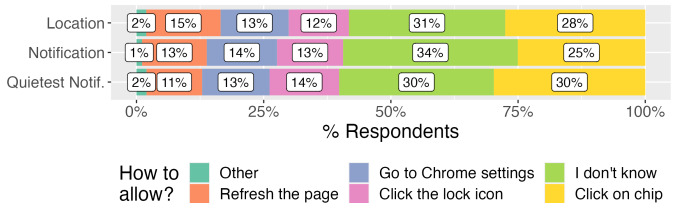

Fig. 9: Respondents' guesses about what they would try first in order to override the quieted prompt.

$\chi^2(30) = 238.3$, $p < .001$, $sresid = 6.7$) as were those with a French locale (42%, $sresid = 2.5$). Respondents with a Spanish locale were more likely to want to refresh the page (18% vs. 13%, $sresid = 3.7$). Respondents taking the survey in Brazilian Portuguese were more likely to want to click on the lock icon (21% vs. 13% overall, $sresid = 6.3$) as were those with Russian locale (17%, $sresid = 3.7$). Finally, those on an English local were more likely to want to click on the chip (31% vs. 27% overall, $sresid = 3.7$).

Across this and the previous findings, we can consider our quieting approach to work if respondents are less than moderately likely to allow after a quieted prompt or identify the correct action to override while noticing the prompt. Based on this, Figure 10 shows that quieting works for the majority of respondents in the situations they encountered. Yet, 32% seeing a quiet chip and 18% seeing the quietest chip may struggle to use the escape hatch when feeling likely to allow. Another 4-8% have a useful intuition for the escape hatch, but may still not find it by themselves as they did not notice the chip itself.

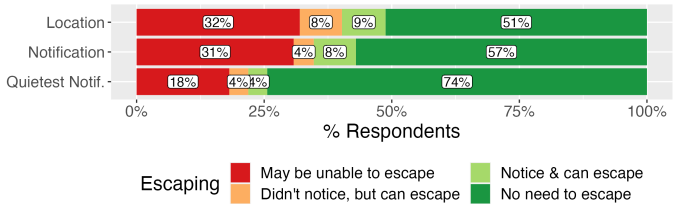

Fig. 10: Distribution of respondents for which we consider our quiet UI to work in the situations they encountered (not at all or only somewhat likely to allow or identifying the correct action to escape) and not work (at least moderately likely to allow and not identifying a useful action) across intervention types.

G. Permanently Disable Quieting (RQ7)

A final aspect of the quieting feature is that users should be able to disable it. Thus, we asked respondents of the second survey what they would first try to disable quieting altogether. Again, respondents were able to choose from a list of five pre-defined actions, including two "useful" options (click the chip, go to settings). Note again that this is primarily based on respondents' intuition, as there has not been any proactive communication about quieting and how to configure it.

As Figure 11 details, 50% indicate that they would try a useful action. There is no statistically significant difference between intervention types ($\chi^2(2) = 10.8$, $p = .005$, all

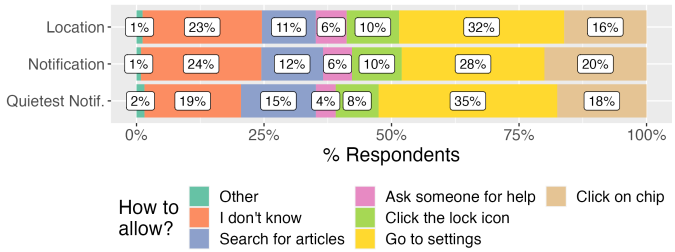

Fig. 11: Respondents' intuition about what action to try to disable quieting entirely.

$sresid < 2$). Similar to the escape hatch, we find that 22% do not know what they should try.

There were again a few statistically significant differences between survey locales. Respondents with a Japanese locale where more likely to report not knowing what they would do (42% vs. 22% overall, $\chi^2(36) = 267.4$, $p < .001$, $sresid = 8.2$). Those using Brazilian Portuguese were again more likely to want to click on the lock icon (16% vs. 9% overall, $sresid = 6.4$). Respondents with a Spanish (8% vs. 5% overall, $sresid = 3.7$) or French (9%, $sresid = 2.0$) locale were more likely to want to ask someone for help. Those using English were slightly more likely to want to go to Chrome settings (34% vs. 32% overall, $sresid = 2.5$).

VII. Limitations

Our study is limited in several ways. For both telemetry and in-product surveys, we are limited to users who did not opt out of telemetry collection. Given that users who opted out made a choice to not share data, it seems at least plausible that they may also exhibit different behaviors and sentiment when it comes to quieting prompts.

Similarly, telemetry collection and in-product surveys are triggered when users visit websites asking for permission. As such, there is an inherent bias towards more frequently visited sites contributing more data points to our results. We believe this is acceptable, as it also represents users' reality, in that they actually encounter such sites more frequently during their day to day of the web. Investigating permission interventions by website or website type can be interesting future work.

For the in-product surveys, imperfect translation as well as differences in general response behavior between cultures may have impacted our results. While we took measures to avoid the former, the latter is hard to eliminate. To make diverging response behaviors transparent, we note statistically significant differences between languages for each analysis we report. Based on this, there does not seem to be an apparent pattern that systematically biased our results. Moreover, due to the short and privacy-preserving nature of Chrome in-product surveys, we do not know how representative the sample we obtained is when it comes to demographic properties such as age or gender. It is at least plausible that not all Chrome users are equally likely to respond to in-product surveys. However, we believe the value of collecting data in context is higher than ensuring perfect representation of the user population.

Finally, there can of course be unwanted and interrupting permission prompts that users end up granting. However, these

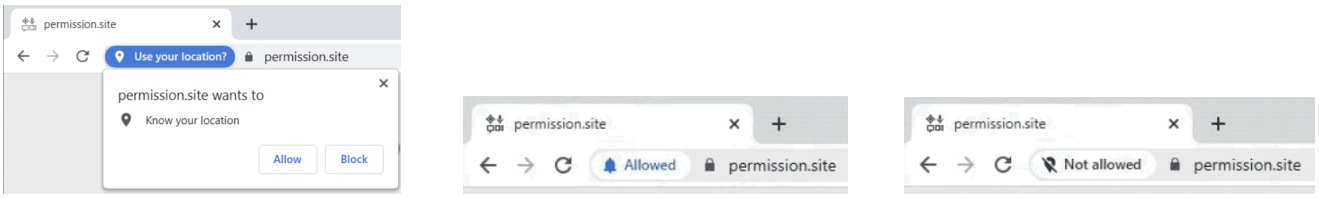

(a) **Request chip**, appearing alongside the permission prompt.

(b) **Confirmation chip** after *allowing* access in a permission prompt.

(c) **Confirmation chip** after *blocking* access in a permission prompt.

Fig. 12: Permission chips introduced in Chrome to reinforce where the permanent permissions control surface can be accessed.

cases are out of scope for this work as they would need a different type of intervention.

VIII. Discussion

Overall, our findings from both telemetry and in-product surveys are encouraging. We discuss the most important aspects in more detail below.

A. Reducing unwanted interruptions and avoiding undesired states

Before WPP's deployment, the previous activation mechanism for the quiet UI [4] was already effective in reducing unwanted interruptions from notifications permission requests. Yet, its impact was limited to 1%-3% of the overall prompt volume for that permission. WPP achieved the same effectiveness (i.e., extremely high precision and recall values) but on a much larger scale. On desktop, it now mediates 43% of notification and 24% of geolocation permission prompts, thus increasing its impact by more than 10x.

Our telemetry data showed that WPP has extremely high post-hoc precision ($> 99\%$) on quiet prompts, despite the ML model being tuned to achieve a pre-hoc precision of $> 95\%$ (on loud prompts). The difference between post- and pre-hoc likely stems from the fact that after the quiet UI is shown (post-hoc), it is less likely to be noticed and thus granted, as compared to the loud prompt.

One potential challenge with showing the quiet UI more often is that it might reduce the amount of prompts with a loud UI, which are used to derive the statistics and features for training the WPP ML model. In order to mitigate such an issue, Chrome clients rely on a *holdback* percentage mechanism to disregard WPP's verdict on a small fraction of prompts that are eligible for WPP's enforcement. The effect is that even in the extreme case where WPP's verdicts resulted in all prompts being eligible for the quiet UI, Chrome clients will still show the loud UI on 15-30% of prompts, which in turn will be used as training samples.

One undesired scenario that WPP wants (and manages) to avoid is for a user to end up in an all-absorbing state, in which only quiet prompts would be shown. That could hypothetically arise when a user starts using Chrome, only dismisses or denies the initial (loud) permission prompts, and WPP's model starts predicting that such a user is very unlikely to grant permission prompts in the future as well. This could result in WPP always showing the quiet prompt to such a user, which is undesirable as users might change behavior or preferences over time. In order to overcome such a potential issue, WPP's 28 days rolling window gradually reduces the count of loud prompts shown to the user to less than 4 (as described in Section IV-A), which in turn deactivates WPP, and makes Chrome resume showing the loud prompt UI again.

B. New quiet prompt UI

The chip treatment appears to fulfill the design goal of being less interrupting than loud permission prompts, as 51% of respondents state to not notice it initially. The quietest chip led to even less respondents noticing it. It is also easy to ignore, as 99% of quiet prompts shown were not overridden. Similarly, 84% of respondents found the intervention helpful and it did not make them feel uneasy. The quietest treatment for sites on which the capability is very unlikely to be granted was received even more favorably. However, we also find room for improvement as detailed below.

C. Providing users with more control

Respondents who felt at least moderately uneasy about the intervention mostly desired more control. While Chrome provides an option to override, this UI appears not to be easy enough to discover intuitively. As detailed in Section VI-F, only 40% of respondents had a useful first guess on what to do to override. While it is likely that, even without a useful initial intuition, others would have been able to override and turn the feature off through trial-and-error, there is an opportunity to improve discoverability and thus reduce feelings of unease.

Additionally, the text currently displayed in the chip (*"Notifications blocked"*) may suggest that Chrome has already made a seemingly permanent decision for the user. The quietest UI treatment – lacking any string – was perceived more favorably, which supports that notion.

In sum, unfamiliarity with the ability to click on the chip as well as an assertive string choice may be leading to a perceived lack of control. To improve this, we started to implement several steps to strengthen user mental models around where permission controls are located in general. For instance, Chrome increasingly leverages chips as a consistent UI pattern for regular permission prompts, in order to associate the location of the chips with where the control surface is. Chrome is in the process of rolling out three kinds of chips:

- **Request chips**, showing a question (*"Use your location?"*) to already start highlighting where the manage-

ment surface is when the permission prompt is showing (Figure 12a);
- **Confirmation chips** (*"Allowed/Not allowed"*), to further strengthen where a decision can be reversed after making a decision (Figure 12b and 12c); and
- **Indicator chips**, showing when a granted capability is actively used by the site and hinting at where it can be turned off.

All these chips show in the same location in the address bar. Clicking on any chip at any time will bring up the site controls surface, which always has permission controls. The request and confirmation chips rolled out in Chrome versions M111 and M109, respectively, while indicator chips were still forthcoming at the time of writing.

While using chips consistently should strengthen discoverability of where to take action, we also plan to change the string in the quieted prompt chip from *"Notifications blocked"* to *"Use notifications?"*, reusing the text of the request chip. While the initial idea with this string choice was to reassure users that Chrome has prevented a site from interrupting them, phrasing this as a question will help to communicate that users can still make a choice. This should provide a heightened sense of actionability and therefore reduce the perceived lack of control. The absence of the prompt itself as well as the text in the popup after clicking on the chip should be sufficient to provide contextual clues about what is happening. Beyond this, we can also consider educational interventions to explain prompt quieting outside of the product.

D. Reduce false positives on sites following best practices

Websites that follow UX best practices when requesting permissions use web capabilities for their intended purpose and provide a clear user benefit from accessing the permission-gated capability. Telemetry data computed on a small sample of 15 popular websites (productivity, news and social media) indicates that users tend to behave differently: even those who frequently deny permissions on other sites are more likely to grant access on these sites. Examples of such sites include messaging sites requesting the notification permission.

Currently, WPP determines which permission prompts to quiet only based on the user's past actions on prompts of the same permission type and on the current browsing context. Hence, as permissions are only granted occasionally, per-user signals carry a lot of weight, often ignoring the fact that some sites may have more popular and helpful use cases for a given permission, resulting in undesired quieting of prompts. We suspect such false positives might also be one of the factors contributing to lower helpfulness ratings and reasons for feeling uneasy reported in the in-context survey.

To further improve the ML accuracy and reduce false positives, WPP could also consider additional features, such as site-related aggregated statistics (e.g., grant/deny/dismiss/ignore rates for a given website) as well as other crawler-based signals.

IX. Conclusion and Next Steps

In this paper, we presented an evaluation of an improved intervention to quiet permission prompts on users' behalf in the wild. Using telemetry, we find that Chrome can now intervene on a substantially larger number of permission prompts, while keeping false positive rates low. This reduces interruptions and prompt blindness even further. In-product survey respondents mostly rate the intervention as helpful without causing substantial feelings of unease. Our results further suggest some room for improvement: the remaining false positives in our evaluation are frequently driven by popular sites, the ability to override Chrome's intervention is not easy enough to discover, and some respondents indicated a lack of perceived control.

We are currently evaluating options to address the shortcomings, as discussed in the previous section. The team will consider a new version of the WPP with improved signals, which could help to reduce false positives that seem to affect some of the sampled sites that follow best practices. Chrome also is in the process of rolling out a consistent chip-based UI along permission prompts to reinforce where permissions can be managed. Additionally, we are planning to change the text in the quiet prompt UI, to more directly invite users to override and thus provide a heightened sense of control.

Acknowledgments

We would like to thank Florian Jacky for his help with fielding the in-context surveys. We are also grateful for Tiff Perumpail, Sabine Borsay, Ceenu George, Mike West, Alisha Alleyne, Nina Taft, Caitlin Sadowski and Adriana Porter Felt as well as the anonymous reviewers helping us to improve the manuscript.

References

[1] Alessandro Acquisti, Idris Adjerid, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Yang Wang, and Shomir Wilson. Nudges (and deceptive patterns) for privacy: Six years later. In *The Routledge Handbook of Privacy and Social Media*, pages 257–269. 2023.

[2] Alistair Dabbs. Enough with the notifications! Focus assist will shut them u... 'But I'm too important!'. https://www.theregister.com/2022/08/05/something_for_the_weekend/, 2022. Last accessed: 2023-06-28.

[3] Pieter Arntz. Browser push notifications: a feature asking to be abused, 2019. https://blog.malwarebytes.com/security-world/technology/2019/01/browser-push-notifications-feature-asking-abused/.

[4] Igor Bilogrevic, Balazs Engedy, Judson L. Porter III, Nina Taft, Kamila Hasanbega, Andrew Paseltiner, Hwi Kyoung Lee, Edward Jung, Meggyn Watkins, PJ McLachlan, and Jason James. "Shhh...be quiet!" Reducing the Unwanted Interruptions of Notification Permission Prompts on Chrome. In *USENIX Security*, 2021. https://www.usenix.org/conference/usenixsecurity21/presentation/bilogrevic.

[5] BlinkOn 15. Day 1 keynote and lightning talks. https://youtu.be/-P_WMKaIhfA?t=828, 2021. Last accessed: 2023-06-28.

[6] Microsoft Edge Blog and Microsoft Edge Team. Reducing distractions with quiet notification requests. https://blogs.windows.com/msedgedev/2020/07/23/reducing-distractions-quiet-notification-requests/, 2020. Last accessed: 2023-06-28.

[7] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W. Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. Your attention please: designing security-decision UIs to make genuine risks harder to ignore. In *Symposium on Usable Privacy and Security (SOUPS)*, 2013. https://dl.acm.org/doi/abs/10.1145/2501604.2501610.

[8] Matt Burgess. Chrome and Firefox are fixing the internet's most annoying problem. https://www.wired.co.uk/article/chrome-firefox-browser-notifications, 2020. Last accessed: 2023-06-28.

[9] Weicheng Cao, Chunqiu Xia, Sai Teja Peddinti, David Lie, Nina Taft, and Lisa Austin. A large scale study of users behaviors, expectations and engagement with android permissions. In *USENIX Security*, 2021. https://www.usenix.org/system/files/sec21-cao-weicheng.pdf.

[10] Google Chrome. Google Chrome privacy notice. https://www.google.com/chrome/privacy/, 2022. Last accessed: 2023-06-28.

[11] Google Chrome. An Update on the lock icon. https://blog.chromium.org/2023/05/an-update-on-lock-icon.html, 2023. Last accessed: 2023-06-28.

[12] Yusra Elbitar, Michael Schilling, Trung Tin Nguyen, Michael Backes, and Sven Bugiel. Explanation beats context: The effect of timing & rationales on users' runtime permission decisions. In *USENIX Security*, 2021. https://www.usenix.org/conference/usenixsecurity21/presentation/elbitar.

[13] Balazs Engedy and Igor Bilogrevic. Permissions misuse & dark patterns. https://www.w3.org/Privacy/permissions-ws-2022/papers/Permission-Misuse-and-Dark-Patterns.pdf, 2022. Last accessed: 2023-06-28.

[14] Adrienne Porter Felt, Alex Ainslie, Robert W. Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettes, Helen Harris, and Jeff Grimes. Improving SSL warnings: comprehension and adherence. In *ACM Conference on Human Factors in Computing Systems (CHI)*, 2015. https://doi.org/10.1145/2702123.2702442.

[15] Marian Harbach, Emanuel Von Zezschwitz, Andreas Fichtner, Alexander De Luca, and Matthew Smith. It's a hard lock life: A field study of smartphone (un) locking behavior and risk perception. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014. https://dl.acm.org/doi/abs/10.5555/3235838.3235857.

[16] @iamdevloper on Twitter. Browsing the web in 2019... https://twitter.com/iamdevloper/status/1090589206013976576, 2019. Last accessed: 2023-06-28.

[17] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Symposium on Usable Privacy and Security (SOUPS)*, 2016. https://dl.acm.org/doi/abs/10.5555/3235895.3235899.

[18] Lloyd Atkinson. Consider disabling browser push notifications on family and friends devices. https://www.lloydatkinson.net/posts/2022/consider-disabling-browser-push-notifications-on-family-and-friends-devices/, 2022.

[19] Nathan Malkin. Contextual integrity, explained: A more usable privacy definition. *IEEE Security & Privacy*, 2022. https://ieeexplore.ieee.org/abstract/document/9990902.

[20] PJ McLachlan. Introducing quieter permission UI for notifications. https://blog.chromium.org/2020/01/introducing-quieter-permission-ui-for.html, 2020. Last accessed: 2023-06-28.

[21] Mozilla. Restricting notification permission prompts in Firefox. https://blog.mozilla.org/futurereleases/2019/11/04/restricting-notification-permission-prompts-in-firefox, 2019. Last accessed: 2023-06-28.

[22] Mozilla. Restricting notification permission prompts in Firefox. https://blog.mozilla.org/futurereleases/2019/11/04/restricting-notification-permission-prompts-in-firefox/, 2021. Last accessed: 2023-06-28.

[23] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.

[24] Patricia A Norberg, Daniel R Horne, and David A Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1):100–126, 2007.

[25] Robert W. Reeder, Adrienne Porter Felt, Sunny Consolvo, Nathan Malkin, Christopher Thompson, and Serge Egelman. An experience sampling study of user reactions to browser warnings in the field. In *ACM Conference on Human Factors in Computing Systems (CHI)*, 2018. https://doi.org/10.1145/3173574.3174086.

[26] Donald Sharpe. Chi-square test is statistically significant: Now what? *Practical Assessment, Research, and Evaluation*, 20(1):8, 2015. https://scholarworks.umass.edu/pare/vol20/iss1/8/.

[27] Karthika Subramani, Xingzi Yuan, Omid Setayeshfar, Phani Vadrevu, Kyu Hyung Lee, and Roberto Perdisci. When push comes to ads: Measuring the rise of (malicious) push advertising. In *ACM Internet Measurement Conference (IMC)*, 2020. https://doi.org/10.1145/3419394.3423631.

[28] Emanuel von Zezschwitz, Serena Chen, and Emily Stark. "It builds trust with the customers" - Exploring user perceptions of the padlock icon in browser UI. In *IEEE Security & Privacy Workshops: SecWeb*, 2022. https://ieeexplore.ieee.org/abstract/document/9833869.

[29] MDN web docs. Introduction to web APIs - Learn web development | MDN. https://developer.mozilla.org/en-US/docs/Learn/JavaScript/Client-side_web_APIs/Introduction, 2020. Last accessed: 2023-06-28.

[30] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. Android permissions remystified: A field study on contextual integrity. In *USENIX Security)*, 2015. https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera.

APPENDIX

A. In-product Survey Questionnaires

$request_type = {"send(ing) you notifications", "see(ing) your location"}

a) Survey 1:

Q0. This website just asked to $request_type. Help us improve how websites ask for permission by taking this 1-minute survey.

Q1. Before seeing this survey, did you notice that this website just asked to $request_type?

- Yes
- No
- I'm not sure

Q2. This website was blocked from $request_type. Why do you think that is? [randomized order]

- This website has a technical issue
- I have told Chrome to block this website
- Chrome thinks that this website is dangerous
- Chrome thinks that I'm not interested in this website
- I previously denied this website's request
- I don't know
- Other (please specify) [not randomized]

Q3. *[if quiet chip]* Chrome automatically blocked this website's request, based on your past choices. How helpful do you find Chrome's action?
[if quietest chip] Chrome automatically blocked this website's request, because most people block it or notifications from this site may be disruptive. How helpful do you find Chrome's action?

- Extremely helpful
- Very helpful
- Moderately helpful
- Somewhat helpful
- Not at all helpful

Q4. How uneasy do you feel about Chrome's action?

- Extremely uneasy
- Very uneasy
- Moderately uneasy
- Somewhat uneasy
- Not at all uneasy

Q5. *[if moderately or more unease, quiet chip]* Please briefly describe what makes you feel uneasy about Chrome blocking requests based on your past choices.

[if moderately or more unease, if quietest chip] Please briefly describe what makes you feel uneasy about Chrome blocking requests that most people block or because notifications from the site may be disruptive.

Q6. Thank you for helping to improve Chrome!

b) Survey 2:

Q0. This website just asked to $request_type. Help us improve how websites ask for permission by taking this 1-minute survey.

Q1. Before seeing this survey, did you notice that this website just asked to $request_type?

- Yes
- No
- I'm not sure

Q2. How likely are you to want to allow this website to $request_type?

- Extremely likely
- Very likely
- Moderately likely
- Somewhat likely
- Not at all likely
- I'm not sure

Q3. Imagine you wanted to allow this website to $request_type right now, what would you try first to do that? [randomized order]

- Refresh the page
- Click on "$request_type blocked"
- Go to Chrome settings
- Click the lock icon
- I don't know what I would do
- Other (please specify) [not randomized]

Q4. If you wanted Chrome to never block a request from a website again, what would you try first to do that? [randomized order]

- Click on "$request_type blocked"
- Go to Chrome settings
- Click the lock icon
- Ask someone for help
- Search for articles describing how to do this
- I don't know what I would do
- Other (please specify) [not randomized]

Q5. Thank you for helping to improve Chrome!