

Privacy Slider: Fine-Grain Privacy Control for Smartphones

FLORIAN BEMMANN, LMU Munich, Germany

HELENA STOLL, LMU Munich, Germany

SVEN MAYER, LMU Munich, Germany

Allow Gmail to access message contents?




Fig. 1. Privacy Slider enables users to select which granularity of their data they want to give to smartphone apps.

Today, users are constrained by binary choices when configuring permissions. These binary choices contrast with the complex data collected, limiting user control and transparency. For instance, weather applications do not need exact user locations when merely inquiring about local weather conditions. We envision sliders to empower users to fine-tune permissions. First, we ran two online surveys (N=123 & N=109) and a workshop (N=5) to develop the initial design of *Privacy Slider*. After the implementation phase, we evaluated our functional prototype using a lab study (N=32). The results show that our slider design for permission control outperforms today’s system concerning all measures, including control and transparency.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: datasets, neural networks, gaze detection, text tagging

1 INTRODUCTION

Ubiquitous and mobile devices use various user behavior data, e.g., mobile behavior [54], location [56], and physiological data [51]. These tracking features enable adaptive and intelligent user interfaces, providing the user with information right when needed, e.g., [37]. For this, users enable various permissions [9] with extensive insights into the user’s personal profile [39]. However, only around 6% of users understand the scope of the permissions they agree to [59]. Shen et al. [59] argued that current mobile systems hardly convey to users what happens with their data and which specific data is used. Moreover, apps request permission, such as location, to display weather information and full device access to support accessibility. However, the weather could be forecasted by only knowing the current city, and the screen readers only need the screen content – not full access. Current permission systems do not allow fine-grain control but only some toggle switches for groups of data access [39]. In combination, users are often unaware of what they agree to and do not have the needed control when understanding the specific case. Thus, users need more transparency in the process coupled with better control of their data.

With smartphones becoming more intelligent, apps and services require increasingly more contextual user data to provide high-quality support, e.g., O’donoghue and Herbert [48]. To reduce potential privacy issues, the researcher has proposed a wide range of mechanisms to preserve the users’ privacy. The simplest solution is outsourcing the decision-making process to an algorithm [22, 46, 52]. However, it takes all control from the user. Thus, Gao et al. [27] opposed recommending decisions to the user and not just taking them for them. Moreover, decisions might be

context-dependent [68]. This raises the question: Are all-or-nothing decisions, as they are implemented with the current grant or deny toggle switches, a good option in the first place? The current operating system addresses this by adding a frequency component to the permission, which the user can set for how long the permission is valid. However, the underlying decision as to which specific data the app has access remains the same. For this, Olejnik et al. [46] added an option “obfuscate” as an alternative to “deny” and “grant.” Here, obfuscate helps to retain privacy which 73% of users found useful. An alternative approach is by Malviya et al. [43], who envisioned that fake data could be used to make a function work while protecting the users’ privacy. While these are options to increase the privacy level of the users’ data, they lack the ability to allow the users to determine the abstraction level of data to be shared. In other words, the user will still be given full device access and meter-precise location. In summary, the permission interface has to meet the tradeoff of giving users detailed control while sufficiently summarizing and reducing the options so that users are not overwhelmed.

In this work, we aim to address the limitations of today’s all-or-nothing permission control systems. As such, we designed and developed a fine-grain permission system. In detail, we envision replacing today’s toggle switch with a *Privacy Slider* to hand full data control back to the user. For this, we first conducted an online survey ($N=123$) to understand what control users envision they need to gain control over their data. In detail, we asked how they could subdivide the all-or-nothing permissions; for instance, the location could be subdivided using the location precision. The results show a total of 135 potential steps for ten datatypes. Next, we investigated the user’s concern concerning the steps using an online study ($N = 109$). This allowed us to rank and potentially group the steps to optimize usability; we supported this step by conducting a workshop ($N = 5$) before starting a multi-stage design and development process of our *Privacy Slider*. In a final evaluation ($N = 32$), we compared *Privacy Slider* to the classical toggle switches in two scenarios: a permission popup and the device permissions screen.

With this paper, we make a set of contributions all leading up to the design and implementation of *Privacy Slider*. Our first two studies show how sliders can support control and transparency based on steps between today’s all-or-nothing choices. The subsequent implementation showed that *Privacy Sliders* significantly outperformed toggle switches with respect to privacy, security, control, transparency, and understandability. Thus, *Privacy Sliders* have the potential to support users in making better decisions when controlling their devices’ permissions.

2 RELATED WORK

In the following, we review the literature with respect to permission systems on smartphones. We first highlight the limitations and drawbacks of current systems. Next, we look at alternative ideas to the current systems, such as automatic systems and fine-grain permission manipulation.

2.1 Limitations of Smartphone Permission Systems

Mobile devices, such as smartphones, collect various data about their users, context, and usage [38]. All major mobile operating systems implement a permission system, where users must grant data access for specific datatypes to apps individually [50]. However, the implemented smartphone privacy concepts face limitations and rarely introduce real privacy from a user perspective [21]. Christin et al. [15] found that existing privacy-enhancing systems lack clarifying privacy implications, and users behave inconsistently with their concerns. Balebako et al. [4] found that users are not careless on that topic, but instead have misconceptions about data sharing that happens through smartphone apps and lack sufficient information. On the other hand, users outweigh anticipated costs and potential benefits, referred to as *privacy calculus* [17]. Here, users accept data being collected in exchange for being able to use a service, c.f. *Price of*

Convenience [35]. Additionally, today we see *digital resignation* [18, 58] or *privacy fatigue* [14] as an overload or lack of control leads to resignation, i.e., users giving up dealing with privacy decisions. As such, users face a challenge with the fundamental concepts of permissions and the associated privacy.

We also see limitations in the user interface itself. For instance, during permission requests, users show low comprehension [23], leading to a *lack of transparency*. Also, apps often do not convey understandably which information leaves the phone, making it hard for users to understand potential privacy leakages [4]. The wording in the permission UI was also found to be hardly understandable, and it was hard for users to grasp the implications [34]. A general *lack of control* is a crucial cause of privacy concerns [42], which has been shown in the online shopping and social media context [65]. Keusch et al. [36] raised concerns about the lack of control. In some cases, users do not even have privacy in their own hands, e.g., if one user leaks a contact list to a service, the other users (who are contained in that contact list) can not do anything against it [50]. The aforementioned two aspects, (1) transparency and (2) control, are identified as the two main pillars of information privacy [6, 28], also coined as the principles of *notice and choice* [55, 71]. A privacy issue introduced by app developers is *permission overclaiming*, also called *permission overdeclaration* [2]. By setting too coarse permissions, developers may claim less data access than they technically have permission-wise. An inappropriately huge amount of permissions also reduces user trust in an application [62]. Fang et al. [21] studied permission overclaiming on the example of the internet permission. This permission poses insufficient expressiveness to enforce control over internet access (i.e., access could be restricted) [5]. They found that many applications would tolerate stricter permission here. 62% request internet permissions, but 36% make requests to specific domains only. Furthermore, third-party libraries that request permission for their purpose lead to that permission being claimed to the full application [49]. Finally, laziness among developers can lead to permission overclaiming due to confusion about the scope of individual permissions [62] and the aim to “just make it work” [5].

2.2 Improving Permission Systems

Fostering **transparency** is one important factor in ensuring users understand the impacts of giving permissions. Thus, a wide range of suggestions to improve transparency has been envisioned, such as adding a purpose to permission toggles Hong et al. [30] (e.g., “file access *for backup*”). Moreover, Cao et al. [12] showed that giving explanations could cut the permission denial rate in their study in half. Mapping permission requests to UI elements helps the understandability [40] about which components of an app actually make use of each permission. Another way to implement transparency in mobile apps is privacy dashboards. In an aggregating or feed-like view, they show the users their collected data. While this method increases transparency, Bemann et al. [6] has shown that users are also getting afraid and feel their privacy being more intruded if transparency is realized without accompanying *control* features.

An opposing approach is relieving the user from privacy decisions through **automatic permission** choice prediction systems, e.g., [22, 27, 46, 52], which try to enhance users’ privacy. However, their effect on actual privacy is disputed; for instance, Elbitar et al. [19] argue that there are no one-fits-all permission decisions. Thus, automatic approaches could be limited in their usefulness. Furthermore, with changing context, users might want to reevaluate their decision, which Wijesekera et al. [68] investigates by detecting such incidents to trigger the permission request again.

Beyond individualized solutions, many technical solutions have been proposed to serve as a middleware between the app and the user as privacy-enhancing technologies (PETs) for smartphones (e.g., [3, 20, 57]), for an in-depth survey see [60]). Moreover, Pennekamp et al. [50] reviewed privacy enforcement strategies on smartphones. On the level of user manipulation, they structure concepts regarding privacy mechanisms into three categories: 1) **reporting** (e.g., omnipresent install prompts, permission visualization, and ways to allow tracking the flow of private information), 2)




Fig. 2. The development process of *Privacy Slider*.

fine-grained tuning (such as user-based configurations), and 3) **fencing** information (e.g., Mockdroid [7], TISSA [72], SHAMDROID [11]).

2.3 Fine-Grain Permission Systems

With the evolution of mobile operating systems, fine-grained control has proliferated in slow steps. Permission popups allowing to choose “only one time” access mitigate the issue of permanent access [44]. Hong et al. [31] already proposed sliders as interface elements as an extension for the one-time-only feature with three options options: allow, ask, and deny. Research proposed various approaches to give users finer control of their data. Jeon et al. [32] categorize permissions into four classes (e.g., outside resources, sensors), each of which common strategies for permission subdivision can be applied to. Zhou et al. [72] enables users to bypass the compulsion to grant a permission to use an application by giving the option to pass falsified data such as empty data, anonymized data, or bogus. Other approaches involve restrictions on how many times a critical resource may be accessed [45], and context-dependent privacy policy configuration [16].

More drastically, Scoccia et al. [57] restructure the Android permission interface by allowing users to (1) make permissions on a feature level and (2) grant finer-grained permissions by introducing permission levels, i.e., a granularity at which data / a resource can be accessed. They found that users appreciated the greater choice, felt more control, and had higher trust. The traditional Android permissions, in contrast, were described as misleading, and the enforced binary choice was not preferred. However, their study is rather proof of the general concept of more granularity in smartphone permissions, emphasizing the realizability in Android. The major part of their contribution is an implemented app instrumenter and its evaluation. The design process of their so-called permission levels has come up rather short.

3 RESEARCH GAP AND DERIVED CONCEPT OF A CONTINUOUS PERMISSION SYSTEM

In summary, the presented related work uncovers technical papers in designing and implementing measures to give control to the users. However, only a few studies investigated the user perspective (e.g., Pennekamp et al. [50] rated usability themselves). Only a few studied finer-grained permissions (e.g., Scoccia et al. [57]) and, if so, emphasize technical aspects rather than the user. Moreover, the review shows that the choices given to the user by today’s systems are typically binary, most prominently toggle switches. At the same time, it is clear that choices are often not binary. For instance, to get the weather forecast, apps do not need the user’s precise location; the city the user is in would be enough to determine whether an umbrella is necessary today. However, current systems only allow an all-or-nothing choice, which causes many privacy concerns and allows for more potential to infringe on users’ privacy than necessary.

As a result, we witness the need for more fine-grained mobile permission systems that enable users to configure their data logging on steps between granting access to all or nothing. For this, we pose three research questions:

RQ1 *What are helpful sub-steps for fine-grained data control?*

RQ2 *How do we deliver the additional control that is usable?*

RQ3 *How does it perform compared to the existing Android permission UI?*

We conducted four studies to address these questions; see Figure 2. In Study I ($N = 123$), we investigated the potential items that are of interest to users to be controllable. In Study II ($N = 109$), we asked users to rate the items concerning their privacy concerns. We used these concerns in Study III ($N = 5$) to rank and group the items into semantically similar groups. Moreover, we investigated possible representation methods, such as how much abstraction and control users want. In combination with multiple rounds of testing and debugging, we developed the final look and feel of *Privacy Slider*. Lastly, we carried out an A/B testing ($N = 32$) to ensure the usability outperforms the industry standard using toggle switches to control users' data.

4 STUDY I: ITEM GATHERING AND CONCERN RATING (ONLINE SURVEY - RQ1)

In this study, we investigate what users would like to control beyond current binary options. Here, we intentionally ask users and experts to gain both perspectives: a) What do users understand to be important for them? and b) What could developers see as important to enable certain applications? Therefore, we conducted an online survey ($N = 123$) with smartphone users and HCI experts to reach a sample of diverse experience levels regarding UX and privacy.

We asked participants to envision the use of 10 different datatypes to explore new avenues for a future permission system. Here, we prompted them with the following ten datatypes: *app usage*, *camera usage*, *incoming message notification*, *phone calls*, *screen content*, *text input*, *user activity*, *voice input*, and *volume & brightness*. The datatypes are rooted in a combination of typically tracked and collected data [54] and common activities [8]. Moreover, all of them can be tracked today and are typically controlled via phone permissions.

4.1 Procedure

First, we explained the procedure and content of the study and asked to give informed consent. Next, we asked participants' demographic data such as age, gender, education, and professional field. See the complete question in the supplementary materials. For each datatype, we collected participants' ideas using the question: "*Which intermediate stages would you find useful?*". Participants did this for ten specific datatypes. Additionally, we asked participants for the logging frequency, which is a property overarching over all datatypes. At the end of the survey, we rewarded participants with 9 GBP per hour.

4.2 Participants

In total, we recruited 123 participants. We recruited 28 participants from our institution and an additional 86 participants via Prolific to diversify the sample. Additionally, we supplemented the sample with 9 HCI experts whom we personally recruited to reflect expert opinions. We required all participants to use a mobile phone or tablet at least almost daily. Participants were between 19 and 74 years old ($M = 29.1$, $SD = 9.4$), and 64 identified as female, 58 as male, and one as diverse. The majority reported a university degree as their highest degree of education (85), 19 had a high school degree, 28 had a high school diploma, 5 had a completed apprenticeship, 4 had a secondary school degree, and one participant finished school without graduation. Their top 5 professional fields were IT, electricians and engineering (45), economy and logistics (14), social and pedagogy (12), services and sales (11), and arts and media (8). In total, our sample resided in 19 different countries, most from Germany (37), South Africa (33), Portugal (10), and the United Kingdom (9).

We determined the targeted sample size on the go via thematic saturation, i.e., when the recruitment of further participants did not reveal new steps [41]. We stopped recruitment when the saturation index reached a threshold of

App usage	Advertisement ¹	App category ⁹	App Name ³⁶	Battery ¹	Duration ²⁹	Frequency ⁸	Location ³¹	Time ²³	User interaction ¹⁸
Camera usage	App name ²	Camera type ¹⁰	Colors ²	Content ²⁵	Date ¹	Duration ²	Editing ²	Focus ¹	
	Frequency ⁶	Lighting ⁶	Location ⁵⁵	Resolution ²	Size ²	Time ²¹	Type ⁹		
Incoming message	App name ⁷	Autoreply ³	Content ²⁵	Deletions ¹	Emojis ¹	Frequency ⁹	Language ¹	Length ⁴	
	Location ³⁶	Participants ²⁴	Read status ⁶	Readability ¹	Sound ³	Time ²⁰	Tone ⁴		
Notification	Content ¹⁴	App name ¹⁷	Frequency ¹²	Length ¹	Location ²⁴	Participants ⁶	Reaction ⁹	Sound ⁶	Time ¹⁶
Phone calls	App name ¹	Date ²	Duration ³²	Frequency ⁴	Location ⁴⁵	Missed Calls ²	Output modality ¹	Participants ⁴³	
	Speaker information ²	Time ²²	Tone ³	Topic ³	Transcript ⁸	User interaction ²	Voice recording ³³	Volume ¹	
Screen content	App Name ⁹	Colors ²	Content ²⁸	Date ²	Duration ³	Frequency ³	Location ²²	Time ⁸	Type ¹³
Text input	App name ²	Autocorrect ⁷	Autofill ¹¹	Generated response ²	Character count ¹	Common used sentences ²			
	Common used word ³	Content ²⁵	Deletions ¹	Emojis ¹	Font ²	Language ²	Length ⁹	Location ⁴⁵	
	Participants ²	Resulting action ¹	Time ⁷	Tone ¹	Topic ¹⁰	Type ⁸	Typing behavior ⁸		
User activity	Activity ⁵⁴	App Name ¹	Battery ¹	Context ³	Duration ¹⁰	Frequency ²	Intensity ²		
	Location ³⁶	Participants ¹	Physical Data ¹⁶	Time ¹⁶					
Voice input	App Name ⁸	Audio clarity ¹	Background sounds ²	Duration ⁴	Frequency ⁴	Language ²	Length ¹	Location ²²	
	Resulting action ²	Time ⁶	Tone ⁴	Topic ⁸	Transcript ⁸	Voice assistant input ³	Voice recording ³⁷	Volume ¹	
Volume & Brightness	Automatic adjustments ⁸	Buttons usage ¹	Duration ⁴	Edit source ⁴	Intensity ²				
	Location ¹	Signal processing ²	Thresholds ⁶	Time ⁶	Volume category ⁴				

Fig. 3. The 135 codes for the ten datatypes that emerged from the 1339 participant statements.

90%, which, through the underlying information weighting model, expresses that the probability of mentions being shared between existing and new participants is 90%.

4.3 Results

We used Python, R, and Atlas.ti to analyze the data and ensure the validity of the responses. We received 1339 individual feedback statements for the different data types. We analyzed our participants' responses using Affinity Diagramming [29]. We formed code groups per datatype. Through this process, we sorted the 1339 statements into 135 distinct codes. On average, each participant contributed 9.9 codes ($SD = 11.9$). Figure 3 gives a visual overview of the final groups. We retrieved the most distinct codes for the datatype *text input* (21 distinct ideas), *phone calls* (16), and voice inputs (16).

The most common mention across many datatypes is the location at which a logging event/data item took place, i.e., the location of a physical activity, the location the user was at when receiving a text message, etc. It was mentioned as the most frequent step in 6 of 10 datatypes and was mentioned second frequently for the remaining 4.

For datatypes where it is appropriate, the name of the respective data item was mentioned often; i.e., for datatype *physical activity* the activity's name, and for *app usage* the app's name. The point of time was within the top 5 steps for all datatypes except *text input* and *voice input*. Also, the duration was mentioned often (e.g., on rank 4 for *activity* and *phone calls*, and rank 3 for *app usage*).

4.4 Summary

In total, our 123 participants came up with 135 distinct codes that are important for the ten scenarios. Many steps are common for multiple data types; for example, *App Name* was mentioned with all but one, or *location* was mentioned for every data type. The frequency with which steps are mentioned varies. While *location* is among the top 3 most mentioned steps for most data types, it is among the two least mentioned steps for data type *Volume & Brightness*. However, we do not know for sure whether the mentioning frequency is an indicator of importance, relevance to the user, or privacy concern. It may rather depend on how present an aspect is in the users' minds. Existing research shows that people are initially rather unaware of privacy risks and do hard naming their concerns unless they are confronted with the topic (c.f. Furini et al. [26]). The order of the collected steps is thus neither given by the survey participants nor naturally by the aspects' characteristics. To obtain a concern ranking, which is required in order to place them on a slider scale, we conducted another study.

5 STUDY II: ITEM CONCERN RATING (ONLINE SURVEY - RQ1)

With the results of Study I, we next investigate how the 135 codes (see Figure 3) are rated with respect to their privacy concern. It will inform the design of potential mechanisms to allow users to make fine-grain console adjustments. Thus, we conducted another online survey ($N = 109$) and let users rate their perceived privacy concerns for the codes.

5.1 Procedure

First, we explained the procedure and content of the study. Afterward, we asked them to consent to the data recording and storage. Next, we asked participants' demographic data such as age, gender, education, and professional field. For each step that resulted from the item gathering study (c.f. Figure 3), participants had to rate their agreement with a statement worded "I am very concerned with my smartphone tracking *duration of the activity (e.g., 1h)*" on a continuous (101 point) slider item [25, 53]. We grouped the items on survey pages by datatype, presenting in total 135 items distributed over 10 pages to the user. Each page started with a short paragraph reminding the participants of their task and context. Steps for the overarching property *frequency* were not ranked, as they are all-time indications that have a natural order. We disclose all study instruments in the supplementary materials of this paper.

5.2 Participants

In total, we recruited 109 participants (42 from our institution and 67 via Prolific). As in the first study, participants had to be fluid in English or German and use a mobile phone or tablet at least almost daily. Their ages are between 18 and 60 years ($M = 29.7$, $SD = 9.5$), with 60 identifying as female, 47 as male, one as a diverse participant, and one who preferred not to disclose. The majority reported a university degree as their highest degree of education (65), 34 had a high school degree, 3 had a secondary school diploma, 5 had a completed apprenticeship, and two participants finished school without graduation. Their top 5 professional fields were IT, electricians and engineering (31), economy and logistics (13), health (11), social and pedagogy (9), unemployed (6), service and sale (5), and arts and media (2). 32 did not identify themselves with the given groups and specified other professions. In total, participants reside in 17 distinct countries. The most represented countries of residence were Germany (43), South Africa (21), Portugal (12), and Greece (4).

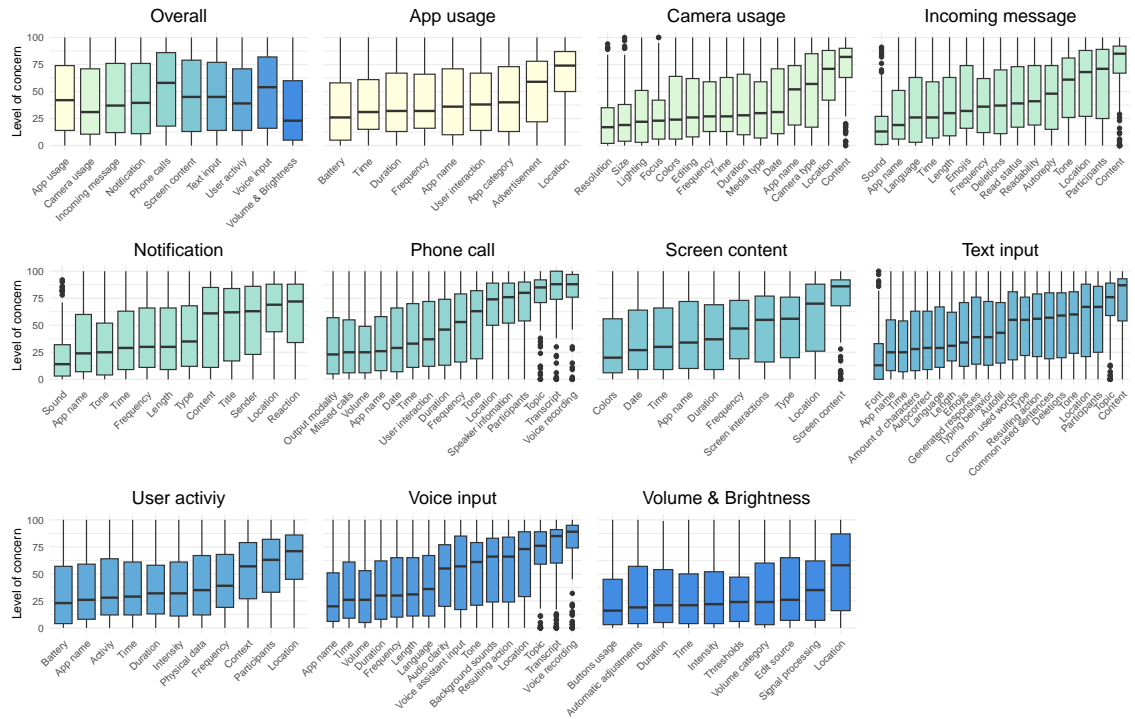


Fig. 4. Boxplot of the concern ratings in our item concern rating study. The black line indicates the median.

5.3 Results

The participant's general privacy concern ratings over all items are in the middle of the 1-100 scale ($M = 45.8$, $SD = 33.8$). The concern values are distributed in a rather bimodal distribution, i.e., the fewest values are in the middle range around 50, and most values are either very low or very high. A slight tendency to the left shows that people were more often rather less concerned than rather high.

Taking a look at the stated concern value grouped by datatype, we see that spoken/written contents were rated most concerning, with phone calls having the highest concern ratings ($Mdn = 59$, $SD = 35.1$) and voice input the second highest ($Mdn = 55$, $SD = 34.8$). Text inputs and screen contents are both on rank three ($Mdn = 46$, $SD = 33.8$). The lowest concern ratings were stated for volume and brightness ($Mdn = 24$, $SD = 31.4$). The values for all ten datatypes are in Figure 4.

5.4 Summary

As a result of study I and II, we created a collection of steps for a set of common data types. We have not performed any reduction or grouping of steps yet. However, to not overload the UI, a reduction to a reasonable amount of steps has to be considered. To group steps, we see two general options: (1) Numerically, i.e. merging steps with close-by concern rating into a group, or (2) semantically. To make this and other design decisions we will in the following conduct a focus group.

6 STUDY III: FOCUSED EXPLORATION (FOCUS GROUP - RQ1)

In the previous two studies, we have collected steps that represent subspects of logged data, accompanied by privacy concern ratings. The results of both studies suggest that it is possible to subdivide today's toggle switches to set up permissions. Study I gave a wide range of steps between all or nothing, and Study II ordered them according to the users' concern level. While this allows us to put all steps on a slider, ensuring full transparency and control; to the users; however, with so many steps, the usability might suffer. Thus, next, we investigate potential user-facing presentations of a novel permission system using a focus group. This investigation will inform the design.

6.1 Procedure

After explaining the focus group and answering any open questions, we asked participants to fill out a consent form and demographics questionnaire. Then the participants introduced themselves and were introduced to the topic. We introduced the general idea and discussed the differences, pros, and cons between toggles and sliders. As the next step, the focus group leader presented the collected steps from Study I augmented using the concerns of Study II for each datatype (see supplementary materials). With this information, we asked the participants to envision to group and potentially sort the steps with the goal of keeping transparency and control high but at the same time also the usability. Finally, the group discussed how the sliders could be best integrated into smartphone usage, i.e., when users would use them and what is needed for effective use.

6.2 Participants

The focus group was conducted with 5 participants (3 female, 2 male) and led by the two first authors. The participants were between 24 and 29 years old ($M = 26.6$, $SD = 2.4$). We were aiming for the perspective of both experts and the user side, so we recruited two experts and three smartphone-experienced users. The experts were HCI researchers who are currently pursuing their Ph.D. The three users were two students and one public service employee.

6.3 Results

The focus group took just over 1 h. We did an audio recording and transcribed this into a text file, which we then coded using ATLAS.ti. This resulted in 49 distinct codes that were assigned in total 90 times. We grouped the codes into 7 code groups, which constitute to topics presented in the following.

Use Case and Target Audience of Privacy Sliders. The focus group identified privacy sliders as **a way to simplify privacy settings for users who are not willing to spend time or don't have sufficient technology understanding** to do so. P1 mentioned *"I don't think the idea of the slider is a bad one because it would make it easier for a lot of people who don't deal with such things that much."* - [P1], and P2 came up with a concrete example of a family member *"[...] I can tell my mom, that once you are halfway through the slider, the app can do more but it also knows more about you."* The participants also pointed out that it is important not to curtail expert users in their control over their data, i.e., still **have detailed control options available on demand** *"But maybe it would be good, as [P2] says, that you would then still have the possibility to set individual things differently."* - [P1] and proposed separate toggles e.g., to turn off single aspects. P2 summarized the combination of privacy sliders and on-demand toggles as *"Good thing for people without much IT knowledge [...], but if I deal with it, then I can use the fine limb ticks."*

Privacy Slider Scope. The focus group came up with the idea of having **one central, default privacy configuration that overarches all apps and datatypes**: "I would expect to be able to set it in the system once, all abide by it and if they need something extra then I am asked." - [P1]. When users are about to perform some action they are least willing to deal with privacy configuration and would benefit from a default configuration "If you want to take a picture, it should be fast but first you have to adjust everything" - [P3]; "Therefore already before!" - [P1].

Finer Granularity for Continuous Datatypes. When the group discussed how much sense sliders would make for specific datatypes, P1 and P2 came up with examples of use cases where a reduced granularity of data would be sufficient. P1: "I have an example for the location theme. If you could then just go in again, and change that again, then it's enough if I can set: City, State, Country. Or I would like to say more precisely that I am currently in XY street." The group had the opinion that ordering steps for a slider specifically on location data is easy, concluding that a **slider to control location granularity** would be good. "Although I would say with location, within locations I would find it exciting if it was a slider. From not at all, to city or urban area at 300m." - [P1]. Similar ideas arose for content, such as texts, speech, and images. "Kind of like the direct content, it is obvious what's on the far right and everything before that is what you can infer through the content. So like tone, emoji, language." - [P2]. A level of "content abstraction" [P2] was proposed as a continuous scale that could be mapped onto a slider.

Grouping of Steps. Participants suggested to rather **group steps by topic** instead of strictly adhering to privacy concern levels. They proposed various groups that make sense to them, e.g. P3 suggesting that "Time and Duration could be put together, so everything that has this time and duration aspect." or P1 who distinguishes between personal and contextual data: "I would try to separate it like this: personal data, the data that is more context and something like location or context plus data related to something like app name". A general **desire for grouping** was expressed especially in cases where many steps exist "It's just a lot. So you couldn't display it like that on the slider, you would have to group it in any case." - [P3].

Ordering of the steps. We discussed the order of steps in the focus group, which was derived from Study II (Item Concern Rating). Participants **overall agreed with the resulting order**, however, the difficulty of deciding on an appropriate order varied with the datatype. P1 and P2 stated that they **did hard ordering the steps of app usage and activity**, while they found that ordering text input went intuitively easy through the degree of content abstraction. The focus group participants could not comprehend why camera type and duration were rated relatively concerning, while physical data and emojis received a surprisingly low concern rating.

The Relation between Data Privacy Concern and Importance. When discussing the privacy concern rating in turn of the step order, the focus group also discussed whether the concern is the right ordering criterion in this case. P3 mentioned that she doesn't find the location very private if it is really necessary: "I don't think that's so bad, because you need the location for many apps. Be it renting a car or scooter or Google Maps.". In line with this, the idea is to order by the ratio of privacy sensitivity and importance in a specific use case.

User Desires and Design Decisions. When a system contains multiple privacy sliders (e.g., for various datatypes), it is important that **consistency of similar steps positions** is guaranteed. Otherwise, users might face unexpected behavior. "The location should always be specified the same. Also, if you now have different sliders and I would now set everything to 60 or so, then I would also expect that it is somehow everywhere the same "safe". And if then suddenly a location is already at 50, then that would be super stupid, just because I do not want to read every time." - [P1]. The exact

position of steps on the slider was not deemed that important, P2 suggested mapping them with equal distances instead of trying to represent the exact concern values.

6.4 Summary

The focus group gave us a good understanding of people’s opinions on how our insights from Study I and II could be fused into a privacy slider design. They agreed on the appropriateness of sliders, especially for datatypes that impose a natural order, such as location or content. A slider interface might especially be beneficial to non-expert users, but we also take away that it is important not to restrict expert users by removing detailed controls. While the idea of having a system-level slider to set a default privacy policy was liked, participants also pointed out the advantages of runtime permissions, aligning with findings in the literature (e.g., Scoccia et al. [57]). The privacy slider design should thus incorporate both concepts.

We noted that, in our focus group, the expert participants had higher speaking shares. We perceived that the non-expert users had hard imagining the slider concept in-depth. We, therefore, conclude that a study with an implemented prototype is necessary to get a sufficient user perspective.

7 PRIVACY SLIDERS: THE FINAL DESIGN (RQ2)

Based on our two surveys, the focus group, and a review of related work, we propose a concept for privacy sliders – a novel user-centered mobile data permission system. Privacy sliders realize two central aspects: First, a simpler, easier user interface that enables quick and easy privacy setting-making. It targets users who are either novices or not willing to spend much time on their data privacy configuration. Second, privacy sliders enable users to choose a custom level of granularity, at which they want to allow to pass data to an app. Both are presented in detail in the following two subsections.

7.1 The System-Level Settings Slider - Sliders as Simplification for Fast and Consistent Privacy Configuration

We propose to implement one slider as a central, default privacy configuration, which overarches all apps and datatypes. A prototype is sketched in (Figure 5): The more to the right the user pushes the system-level slider, the more detail is granted for every datatype. To meet expert users’ needs and individual needs on specific datatypes, the access level to a datatype can be overwritten. One slider per datatype allows overriding the system-level slider’s setting (e.g., in Figure 5: For the location, the user has configured lower granularity data access).

Based on the results of our focus group we envision that this simplifies users indicating their privacy preferences. Especially novice users and those who do not want to spend much time on privacy configurations might benefit from the intuitive and fast UI of a slider.

7.2 Enhanced Permission Popups: Information Minimization of Continuous Datatypes

In the focus group, we found that for some data types, such as location and content (text, speech, and camera were mentioned), it makes sense to configure granularities. Steps for location data could, for example, be reduced to an accuracy of +/- 500 m, city, or country. For many use cases, that might be sufficient: For example, when using a weather forecast app it would be sufficient if the OS passes the city name to the app, instead of the user’s precise location. For content, such as text messages, content abstraction procedures could similarly be applied.

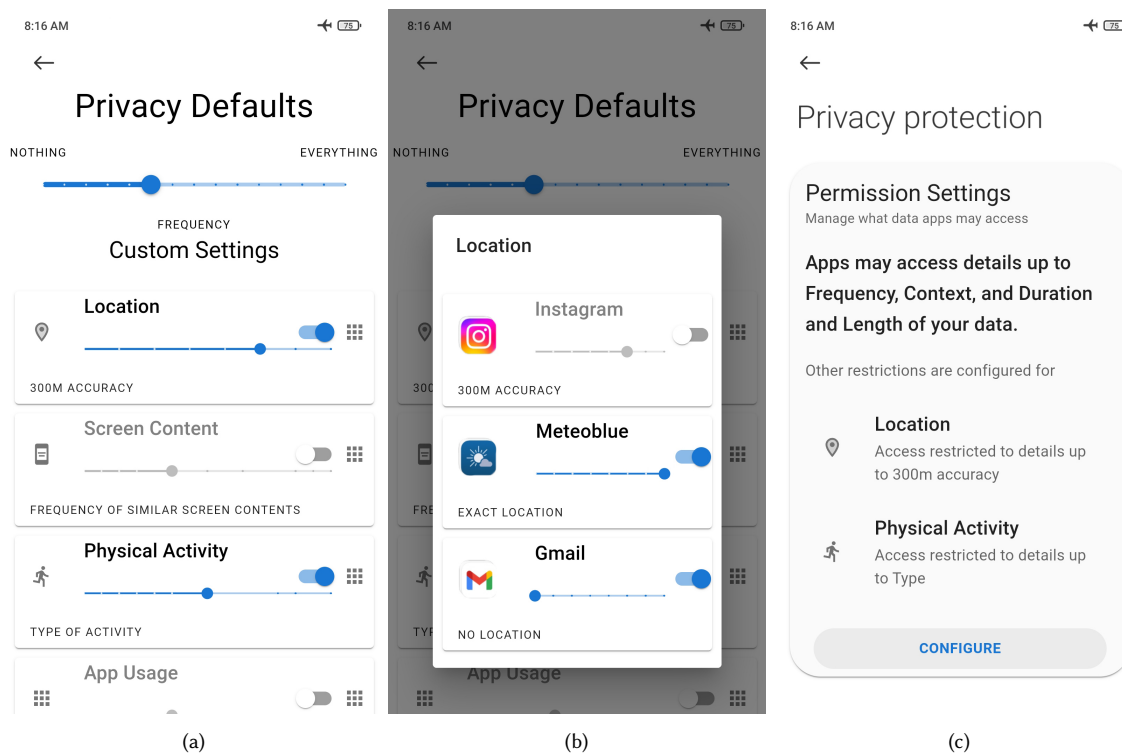


Fig. 5. The system-level slider (a) is used to configure the phone’s general privacy settings, per default applying to all apps and overarching all datatypes. It is especially targeting users who do not want to spend much time with single privacy decisions. The sub sliders below allow to set overriding configurations for single datatypes. A popup (b) allows to make settings per app, and the overview screen (c) summarizes the settings.

Supplementing the previously presented system-level slider, our privacy slider concept introduces such configuration options for location and content. These sliders (see Figure 6 (a) and (b)) are meant to replace the current permission UIs, e.g., the *only one time, always, when using the app* single-choice radio button interface that Android currently uses for location access permission.

8 STUDY IV: SLIDER VALIDATION (LAB STUDY - RQ3)

To evaluate our privacy slider concept, we implemented it as a prototype and conducted a lab study. Participants were asked to use both a UI mockup of a traditional permission interface and a mockup of the privacy slider interface concept. We assessed both interfaces’ effects in a mixed-method approach, using survey items and interview questions. The study consisted of 4 scenarios, with participants going through them two times, once using traditional Android permission UI and once more using the slider interface. Three of the scenarios were runtime permission popup situations (see Figure 7), and the remaining one was the general privacy settings menu deploying our system-level slider (see Figure 5). The order in which the four scenarios were presented was randomized.

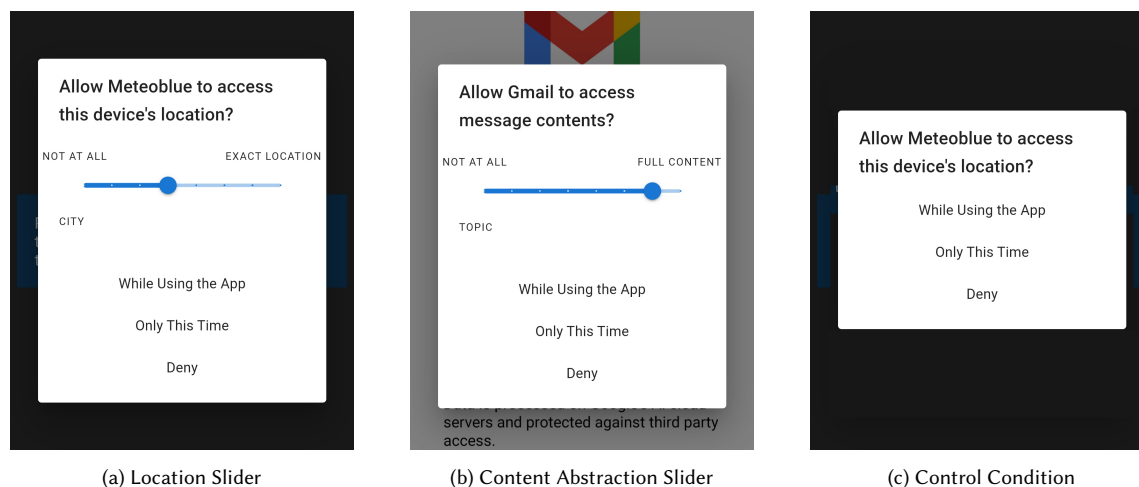


Fig. 6. We enhance permission popups with a slider that allows to choose a level of granularity. (a) visualizes this on the example of the location permission, (b) for text message contents. On the very right (c) we show our control condition, consisting of a slider-less permission popup as it is implemented in the Android UI nowadays.

8.1 Apparatus

We mocked the Android permission UI (runtime popups see Figure 6 (c), and the settings menu see Figure 5 (c)) with a Progressive Web App¹. The runtime permission popup scenarios consisted mainly of a series of app screenshots that the users clicked through, augmented with button respectively slider UIs that mimic the permission interface. Launched in fullscreen mode and being tailored specifically to the study device, the UI experience was very close to real Android UI. To rule out the effects of the mock, we also mocked the traditional Android permission popup UI with this approach, instead of using the OS implementation.

8.1.1 Slider Steps: System-Level Slider. In Section 4 we collected potential steps for privacy sliders for each datatype and ranked them by their level of privacy concern in Section 5. However, in Section 6 we found that a strict order by privacy concern does not make sense to users, as the varying order of similar steps on different datatypes might lead to unexpected configurations. We thus follow the idea of the focus group to group the steps by topics "that make sense". These groups then constitute the steps of the system-level slider and its sub-sliders (except the continuous datatypes location and content). Thus the sliders for all datatypes are designed equal. We order the steps on their slider by the median concern value that our participants in Study II rated them.

8.1.2 Slider Steps: Location and Content. As pointed out by the focus group, location and content pose an inherent granularity, which can be mapped to a continuous slider design. We chose the following steps, based on mentions from the focus group and proposed order in Section 5:

Location: not at all, country, state, city, urban area, 500m, 300m, street name, exact location

Content: not at all, language, length, tonus, emojis, common words/sentences, topic, raw content

¹<https://web.dev/progressive-web-apps/>, last accessed 2024-04-29

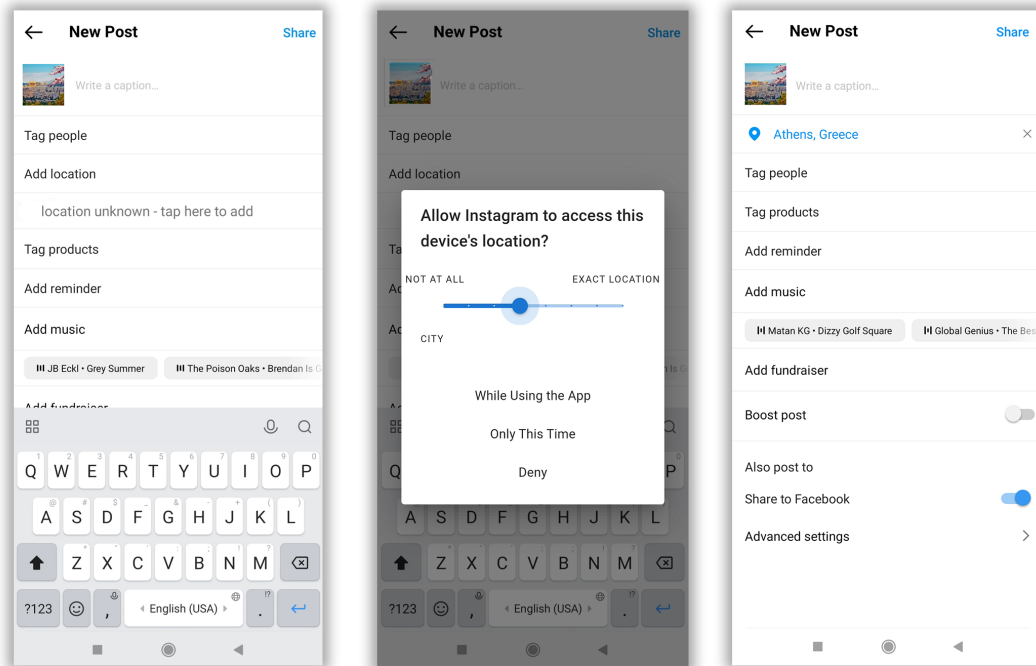


Fig. 7. A series of screenshots that shows one of the scenarios that we used in our studies from left to right: Here the participant is advised to craft an Instagram post, that is tagged with its location. This figure shows the privacy slider condition of the experiment.

8.2 Procedure

The study conductor met each participant in our lab in a separate room with a table. After explaining the study, the participants read and signed the consent form. We then started with a questionnaire on one's individual information privacy concern level using the UIIPC questionnaire [42]. Furthermore, we assessed affinity for technology interaction (ATI) [24] and demographics.

To get participants into the thinking mode and to affiliate with the scenario of giving permissions, the study conductor talked with them for a couple of minutes about their last contact with smartphone permissions, what they thought and felt in that situation, and what the decision was like. Then we introduced them to one of the two conditions *Classic* or *Slider* in randomized order. For both conditions, we ran participants through the same procedure: a demo of the condition, then they tested the three showcases, and finally, they tested the system-level settings application. After each but the demo, they filled in a system usability scale (SUS) [33] questionnaire and answered six items on perceived control, privacy, security, making sense, transparency, and understanding. These items were assessed on a continuous slider scale, we disclose the wording in the supplementary materials. At any time, participants could verbally or in writing articulate additional feedback.

In the end, we asked participants additionally if they had any further feedback. We audio-recorded the full procedure and rewarded participants with 10 EUR per hour or the respective amount of study credit points. Participation took approximately 30 minutes to complete the study.

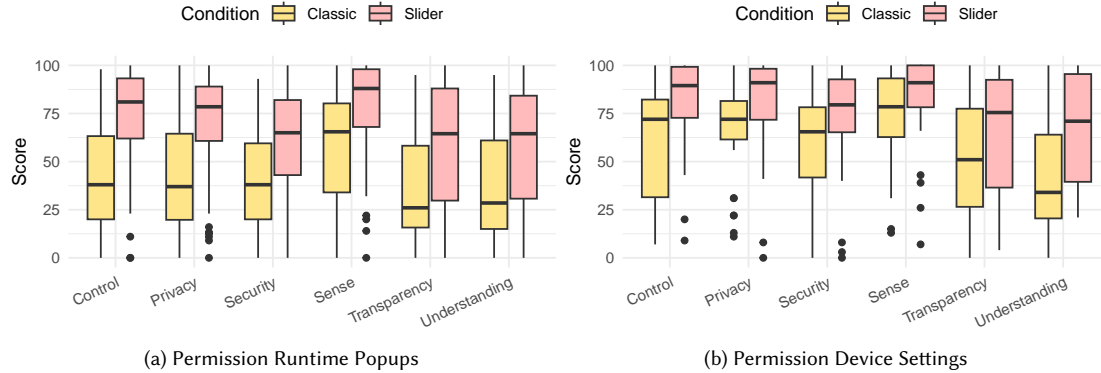


Fig. 8. Ratings on five aspects around privacy compared for the classic permission UI and the slider UI. On the left (a) regarding permission popups on runtime, on right side (b) for the device’s settings menu.

8.3 Participants

We recruited 32 participants via our university mailing list, Slack channel, Instagram, and personal contacts. We required participants to be daily smartphone users and be fluent in English. Participants were between 21 and 70 years ($M = 28.0$, $SD = 8.5$), with 18 female and 14 male participants. They reported having a Master’s degree (16), a Bachelor’s degree (9), a high school degree (4), a doctoral degree (1), not finished school (1), and a vocational education (1). All participants reside in Germany, besides one participant from the United States. To understand our sample’s privacy perception, we assessed the IUIPC questionnaire [42]. Our participants rated their Awareness on average with 6.2 ($SD = 1.3$), Control with 5.7 ($SD = 1.3$), and Collection with 5.7 ($SD = 1.2$) (higher scores mean more privacy-affine). Their mean score of affinity for technology interaction (ATI) was 4.1 ($SD = 1.0$). This indicates a rather technology-affine sample. According to the classification of Franke et al. [24] the ATI of an average population is to be expected at around 3.5, with high ATI samples around 4.

8.4 Results

We run the statistical evaluation in Python and R. Moreover, we applied non-parametric tests when the normality was violated. We did the qualitative coding again in Atlas.ti.

8.4.1 Usability. The users’ rating on the system usability scores (SUS) was higher for the slider UI than for the classical UI, for both the runtime permission popup comparisons and the system-level settings menu, see Table 1. According to the adjective classification of Brooke [10] all interfaces’ usability can be regarded as *excellent*, only the classic device settings menu was rated as *good* only. However, only in the runtime comparison does the slider significantly outperform the classic approach, see Table 1.

8.4.2 Privacy Effects. For the system-level settings, we can use the classical Wilcoxon signed-rank test after confirming the non-normality of the data. However, for the permission runtime popups, we perform an ART-ANOVA [70] with task and participant as random factors to account for the differences between the three showcases. We compared the effects of the classic and the slider UI on their users’ privacy perception, see Figure 8. We show that *Privacy Slider* outperforms the classic approach significantly in nearly all measures. The only non-significant items are the measures

for usability and making-sense in the system-level settings menu. However, descriptively the slider performs better also for these items. See Table 1 for all measures and test results.

8.4.3 Qualitative Feedback. We coded the free text responses, transcribed audio recordings, and interview notes in Atlas.ti. We then organized the codes into code groups, which constitute the following topics.

Slider Interface is Preferred over Classic Button Interface. In general, comments on the slider interface were better than on the classic button-only interface. Many participants mentioned that they did prefer the slider interface, while none said that they’d rather like to stay with the classical version. P39 described it as “a big improvement over the usual UI and would definitely prefer this in all cases.”. Similarly, P12 “No, I think its a great addition.” and P21 “liked it much better than the previous one.”. In contrast, the classic condition was described as having “not enough options for data privacy” (P26) and being “too general”. P53 expected a “more detailed option display”. For continuous datatypes, such as location, the slider interface was more intuitive to use for some participants. One mentioned that they’d prefer it for continuous datatypes only: “Slider for me is only useful for continuous values like distance.” (P51). Especially in the weather app scenario the ability to configure a granularity makes sense to the participants, as e.g. P60 stated: “The weather apps don’t require my exact location so I like this slider feature here”.

System-Level Settings: Better Overview and Easier Getting-Into with Sliders. People liked the ability to use sliders for the device-wide privacy settings as well. Participants liked the overview that the sliders gave on data collection. P39 reported on the system-level slider: “This was close to perfect, I think this is how it should be. The granularity of specific sliders also grants insight into all the different data that is being collected, which the normal UI completely lacks.”. In contrast, the classical settings menu is often criticized as it is hard to get into it and lacks an overview. “Especially at first glimpse, the system is not that easy to get an overview of.” (P62). The lack of transparency of the classical interfaces also leads to a lack of trust, as a further mention of P39 shows: “Many options were hidden and you have to kind of guess what each thing does. Also, not sure if the options will be reverted after an update.”. P26 complained that “you have to click through more” with the classic interface. Besides clear benefits, our participants also pointed out some drawbacks of the new slider interface. Criticism mainly evolved around the system-level slider, which stood on top of the settings screen above all individual permission sliders. E.g. P22: “I don’t think I would use one slider on the top. Especially as the different categories only make sense for some apps. Instead, I would get rid of the slider on top and instead have a categorical setting like privacy level high/medium/low for example.”. P10 saw privacy risks for lazy, speeding users introduced by the system-level slider: “On further reflection, this feature now strikes me as very risky to dangerous. For example, it could

Table 1. The statistical results of Study IV. * we report F values for all but SUS using the ART-ANOVA.

	Runtime Popups								Device Settings							
	Classic		Slider		Normality		Wilcoxon*		Classic		Slider		Normality		Wilcoxon	
	M	SD	M	SD	W	p	W/F	p	M	SD	M	SD	W	p	W	p
SUS	81.9	9.2	88.2	11.1	.925	<.001	88.5	<.001	78.0	17.0	80.0	18.1	.887	<.001	203	.382
Control	39.7	26.2	73.7	24.8	.939	<.001	133.98	<.001	60.5	27.7	80.7	23.6	.876	<.001	63	<.001
Privacy	40.6	25.9	70.5	25.0	.939	<.001	107.21	<.001	66.1	24.6	80.4	25.2	.86	<.001	83.5	<.001
Security	39.8	24.4	61.9	26.1	.964	<.001	65.951	<.001	59.8	25.9	73.5	27.2	.9	<.001	84.5	<.003
Sense	59.7	28.2	81.0	21.2	.885	<.001	52.23	<.001	73.2	25.0	82.2	23.3	.835	<.001	117	.086
Transparency	33.6	24.9	58.6	31.3	.93	<.001	54.58	<.001	48.8	31.5	66.7	29.9	.922	<.001	114	<.009
Understanding	35.2	26.2	58.2	29.4	.935	<.001	54.183	<.001	40.6	29.6	66.2	29.4	.92	<.001	58.5	<.001

tempt me as an annoyed user to be happy to easily set the location permission for all my navigation and sports tracking apps to maximum, and thus unintentionally set e.g. the memory access permission for all apps ever downloaded to completely open as well."

Besides some points of criticism, the participants overall liked the slider version. *"It is very sexy, please install it on every phone."* (P40). They saw the main benefit of the slider-based system-level settings menu in its intuitive understandability and overview. *"The permission settings were very clear and concise, it was easy to gain an impression of how the data would be used."* (P53). Less technology-experienced users would benefit: *"[The slider] controller [is] more intuitive for older people or people who do not have smartphone affinity."* (P25).

Sliders Improve Transparency. Besides improved control, participants also perceived higher transparency about the data collection. The sliders with their steps make transparent which aspects a permission encompasses *"The granularity of specific sliders also grants insight into all the different data that is being collected, which the normal UI completely lacks."* (P39), and even give the user more sense about how their data is used *"The permission settings were very clear and concise, and it was easy to gain an impression of how the data would be used."* (P53). Having an overview of the active steps of each permission slider, the user could quickly grasp what is collected *"They explained what exactly would be collected."* (P22). P28 further saw an explanatory effect and triggered reflection processes: *"Offers control but also explains the usage of the data, and by showing the different levels of data abstraction people get a feeling of how much the data can actually capture and gives an opportunity to realistically reflect on their own boundaries."*

More Perceived Control on How Data is Used. Participants mentioned that the system-level slider gives them control on how data is used: *"The permission settings were very clear and concise, and it was easy to gain an impression of how the data would be used."* (P53). However, we found that, independently of the applied method, participants felt a general lack of control over what happens with their data after granting access to it. Especially regarding the classic interface, many participants mentioned that the given control options only give control on what data is *passed* to an app, but not at all what thereafter *happens* with their data; i.e. with whom it is shared, how it is processed, where it is stored. We observed a general lack of trust in all that happens behind/after the granting interface, independent of the method. It was mentioned that *"trust [is] missing, the method is not the problem"* (P13). Similarly, P38 expresses issues with control over the later stages in the processing pipeline: *"This Interface suggests some form of privacy control but unless one denies everything, there is limited control once data is in the app."* P28 expressed missing *"control over where the data is stored, with whom it is shared, and what information is drawn from it."* This issue is out of the scope of permission granting methods which we focus on in this paper, but nevertheless noteworthy for future work.

Slider Design. A couple of participants expressed different preferences of the slider's step order, for example, P21 generally expressed that *"The order of some levels of privacy didn't make much sense to me."* or P13 proposing based on their cultural background that *"emojis should go to the last. In India, different emojis are differently interpreted."* Customization of steps per app was suggested by P38 *"The slider might be adjusted by application, since for example weather is no more accurate than a couple of 100 meters anyway."* On the other hand, some participants were concerned about too many differences between the sliders. P10 said that inconsistencies in the slider steps could lead to unexpected behavior. A medium slider value should express a similar level of data and privacy across all sliders. In general, it was perceived as a *"very sufficient interface, and with a bit of background knowledge, it is easy to understand how it works and what it does."* (P52).

Detailed UI Comments. As for the nature of a high-fidelity prototype study, participants also pointed out many detailed UX improvement suggestions and criticisms of our prototype. During the study, a couple of unclarities in the UI were pointed out, especially regarding the slider-based system-level settings (e.g. “*not clear what gradations mean.*” (P51)). The behavior of the sub sliders and their toggles was unclear to some participants (“*unintuitive what you turn on with the toggle?*” (P10)), also explanations on the slider steps, for example with context menus, were desired. A few participants generally misconceived the slider steps as selecting instead of summing up. However in general the participants made themselves familiar with the slider UIs quickly, and further explanations by the study conductor were necessary in individual cases only.

8.5 Summary

Summing up on the lab study where we tested our high-fidelity privacy slider prototype with 32 participants, we conclude that the concept of privacy sliders was perceived very positively. Participants saw benefits in several scenarios, the usability was rated better than with the traditional privacy settings UI, and we found positive effects on transparency and control. We collected points of criticism that can be improved in future iterations. For example, step positions on the sliders should be determined by their concern level instead of the percentage of the slider scale, and the cumulative behavior has to be better explained.

9 GENERAL DISCUSSION

9.1 Runtime Permission Sliders Outperform the Standard Permission UI

The feedback on runtime permission popup sliders was overall very good. Extending the current button UI with a continuous choice of data granularity made sense to our participants and was perceived as intuitive, regarding the system usability score it significantly outperformed Android’s current UI. Especially for data types that impose a natural degree of granularity (such as location), it was liked, and participants envisioned situations where they see an advantage in continuous permissions. With its straightforward user flow, which is close to Android’s current design, users got into it easily and there is not much that could trigger confusion. We argue for including this in future runtime permission popups. While fine-granular permission concepts have been published in the past occasionally, for example by Jeon et al. [32] and Scoccia et al. [57], the present study is, to the best of our knowledge, the first study that implements it as a well-usable slider UI and studies its usability and applicability with lay smartphone users.

9.2 System-Level Settings Slider

Also in the system-level settings menu sliders were preferred by our participants. We see a benefit, especially in the transparency and overview that they provide, what participants confirmed in their qualitative statements. The system-level slider on top was deemed a good feature for novice users, and the flat menu structure required users to click through less. However, the more complex nature of a whole settings menu in contrast to a single-case runtime popup makes designing challenging. This reflects in a couple of remaining usability issues that we found in our study, and need to be addressed before rolling this out in the wild.

Most importantly, user support should be included in helping users get into the principle of how the slider-based menu is working, such as a tour, as proposed by Carlèn [13]. Furthermore, it has to be ensured that users do not experience unexpected behavior across the sliders. In our study, we found that users perceived privacy concerns of

specific steps differently, and thus would have expected a different order. However, with the system-level slider on top, it is essential that the subsliders that are moving alongside do not show unexpected configurations.

9.3 A Method-Independent Lack of Control of What Happens with the Data

In the qualitative feedback, participants mentioned that they desire more transparency and control over what happens with their data after a permission has been granted. As they also admit, this is out of the scope of our study on the permission granting interface, however, we think this finding is nevertheless important to note. Sliders could for this issue be part of the solution as well. In our case of granting data access, sliders enabled control and conveyed transparency to their users. Generalizing privacy sliders to a modality for configuring data transactions, they could also find application on other stages of the data pipeline. The setting options of how far data is passed on, or in what depth it is analyzed, pose a natural order (for example data not leaving the device, going to the app company’s server only, being processed on a cloud server, being disclosed to third party companies, being made available publicly).

9.4 The Tradeoff Between Warning Fatigue and User Control

Permission interfaces have to deal with the tradeoff of warning fatigue, i.e. users’ desire for control and on the other hand being mad about too much information, options, and time spent. Users tend to ignore privacy-enhancing technologies (also coined *the challenge of user ignorance*) [1, 4] and concepts that foster their usage have to be considered, such as nudging approaches (e.g. Thompson et al. [63]). Control-providing concepts have to be designed with care, as sophisticated concepts may quickly annoy their users and thereby fail [50]. While this is not the case for our runtime permission slider, users don’t have to deal with the system-level settings unless they actively look for them in the settings menu. Users could be triggered to use the system-level settings in appropriate situations. That could be once after the device has been set up, to choose the device’s default privacy policy, after the installation of multiple new apps, or after context changes (for example when travels or holidays are detected).

9.5 Contextual Privacy and Personalization

The privacy slider configuration could also be context-dependent. Users prefer data disclosure differently depending on their context, as Wang et al. [66] show in the example of online behaviors. Including context in runtime permissions help users make their decision [57] and enables an even more fine-grained choice, also called *flexible permission*[57]. A system that is able to understand and extract a contextual difference given, that would even be possible without additional user burden. Contextual privacy has shown to be beneficial in various use cases, such as online privacy policies [69] or IoT [47]. Regarding smartphone permissions, research has shown that the incorporation of contextual cues can improve decisions [67] and studied machine-learning-based decision support [64]. We think that our fine-granular approach to permissions well integrates with such approaches. The non-binarity of continuous privacy configurations could be used to reflect model uncertainties, i.e. instead of a prediction model requiring to output a binary all-or-nothing decision, an insecure prediction could lead to a slider value somewhere in the middle. Furthermore, the context could be another dimension of configuration that could be controlled through a slider (e.g. rating private situations on weekends as more concerning and worthy of protection than behavior during office days).

9.6 Privacy Sliders Enable Novel Adaptive Use Cases

To avoid the unpleasant consequences of privacy issues, data usage by applications is restricted. Access to potentially sensitive resources, such as screen contents and detailed device activity, is in Android for example organized into

the Android Accessibility services. Thereby access is highly restricted to a few purposes only. By going from the current binary approach to fine-grained configuration we envision that such resources could be opened to wider application purposes. Screen contents for example could be leveraged for adaptive application scenarios, such as predicting next-action sequences, if the data was abstracted to the smallest necessary level of detail. Exact text contents, like text messages, names, or login credentials could be abstracted to tokens like *textmessage*, *name*, and *logindata*. They would thereby still be useful for several application scenarios but way less privacy-invading. Continuous permissions, realized through privacy sliders, thus not only have a privacy-preserving effect on the users but also enable novel opportunities for application and system developers.

9.7 Future Work: Field Study

In the present study, we have focused on users' opinions on the privacy slider. We deliberately asked end-users because they are the major stakeholders regarding privacy; their data is worked with and they opt for buying and using their smartphone. In the next iteration of the privacy slider, developers should also be taken into account, to see which effects fine-granular permissions would have on them, what changes from the development perspective, and find solutions on how developers could deal with that. Developers would need to deal with data of different granularity. If they in the worst case simply reject all except the finest data levels and thus force users to push the privacy slider to the finest level, not much would be won for the user. They still would be in the dilemma of granting (full) data access or not using the application (c.f. Stach and Mitschang [61]). Flexible data structures might be needed in mobile app frameworks, to make it easy to work with varying granularities of data. Furthermore, an in-the-wild study has to be conducted. Our lab study was appropriate to get first insights on privacy sliders, showed that it is promising, and yielded valuable insights for the next iterations. However, to see how users actually use them in real situations and which effects that has, a field study, where privacy sliders are distributed to the users' own devices, has to be conducted.

10 CONCLUSION

In this paper, we designed and evaluated *Privacy Slider*. They go beyond the current binary permission decisions and thereby empower users to make fine-grain privacy decisions. We ran two online surveys (N=123 & N=109) and a workshop (N=5) to develop the initial design of *Privacy Slider*. A lab study (N=32) that we conducted based on an implemented prototype showed that users prefer privacy sliders over the current smartphone permission interface. They are especially advantageous for novice users and when applied in runtime permission decisions. Furthermore, the privacy slider outperformed the classic interfaces in all measures, including increased perceived control and transparency.

REFERENCES

- [1] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (*CHI '15*). Association for Computing Machinery, New York, NY, USA, 787–796. <https://doi.org/10.1145/2702123.2702210>
- [2] Kathy Wain Yee Au, Yi Fan Zhou, Zhen Huang, Phillipa Gill, and David Lie. 2011. Short paper: a look at smartphone permission models. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices*. 63–68.
- [3] Michael Backes, Sebastian Gerling, Christian Hammer, Matteo Maffei, and Philipp von Styp-Rekowsky. 2013. AppGuard-Enforcing User Requirements on Android Apps. In *TACAS*, Vol. 13. Springer, 543–548. https://doi.org/10.1007/978-3-642-36742-7_39
- [4] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. 2013. "Little Brothers Watching You": Raising Awareness of Data Leaks on Smartphones. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (Newcastle, United Kingdom) (*SOUPS '13*). Association for Computing Machinery, New York, NY, USA, Article 12, 11 pages. <https://doi.org/10.1145/2501604.2501616>

- [5] David Barrera, H Güneş Kayacik, Paul C Van Oorschot, and Anil Somayaji. 2010. A methodology for empirical analysis of permission-based security models and its application to android. In *Proceedings of the 17th ACM conference on Computer and communications security*. 73–84.
- [6] Florian Bemmam, Maximiliane Windl, Jonas Erbe, Sven Mayer, and Heinrich Hussmann. 2022. The Influence of Transparency and Control on the Willingness of Data Sharing in Adaptive Mobile Apps. *Proc. ACM Hum.-Comput. Interact.* MobileHCI (2022), Issue 6. <https://doi.org/10.1145/3546724>
- [7] Alastair R Beresford, Andrew Rice, Nicholas Skehin, and Ripduman Sohan. 2011. Mockdroid: trading privacy for application functionality on smartphones. In *Proceedings of the 12th workshop on mobile computing systems and applications*. 49–54.
- [8] Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger, and Gernot Bauer. 2011. Falling Asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (Stockholm, Sweden) (*MobileHCI '11*). Association for Computing Machinery, New York, NY, USA, 47–56. <https://doi.org/10.1145/2037373.2037383>
- [9] Bram Bonné, Sai Teja Peddinti, Igor Bilogrevic, and Nina Taft. 2017. Exploring decision making with Android’s runtime permission dialogs using in-context surveys. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 195–210. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/bonne>
- [10] John Brooke. 1996. Sus: a “quick and dirty” usability. *Usability evaluation in industry* 189, 3 (1996), 189–194.
- [11] Lucas Brutschy, Pietro Ferrara, Omer Tripp, and Marco Pistoia. 2015. Shamdroid: gracefully degrading functionality in the presence of limited resource access. *ACM SIGPLAN Notices* 50, 10 (2015), 316–331.
- [12] Weicheng Cao, Chunqiu Xia, Sai Teja Peddinti, David Lie, Nina Taft, and Lisa M. Austin. 2021. A Large Scale Study of User Behavior, Expectations and Engagement with Android Permissions. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 803–820. <https://www.usenix.org/conference/usenixsecurity21/presentation/cao-weicheng>
- [13] Filip Carlén. 2017. User Onboarding An investigation in how to increase the activation of new customers using design. (2017).
- [14] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. 2018. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior* 81 (2018), 42–51. <https://doi.org/10.1016/j.chb.2017.12.001>
- [15] Delphine Christin, Andreas Reinhardt, Salil S Kanhere, and Matthias Hollick. 2011. A survey on privacy in mobile participatory sensing applications. *Journal of systems and software* 84, 11 (2011), 1928–1946. <https://doi.org/10.1016/j.jss.2011.06.073>
- [16] Mauro Conti, Vu Thien Nga Nguyen, and Bruno Crispo. 2011. Crepe: Context-related policy enforcement for android. In *Information Security: 13th International Conference, ISC 2010, Boca Raton, FL, USA, October 25–28, 2010, Revised Selected Papers 13*. Springer, 331–345.
- [17] Mary J Culnan and Pamela K Armstrong. 1999. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization science* 10, 1 (1999), 104–115. <https://doi.org/10.1287/orsc.10.1.104>
- [18] Nora A Draper and Joseph Turow. 2019. The corporate cultivation of digital resignation. *New media & society* 21, 8 (2019), 1824–1839. <https://doi.org/10.1177/1461444819833331>
- [19] Yusra Elbitar, Michael Schilling, Trung Tin Nguyen, Michael Backes, and Sven Bugiel. 2021. Explanation Beats Context: The Effect of Timing & Rationales on Users’ Runtime Permission Decisions. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 785–802. <https://www.usenix.org/conference/usenixsecurity21/presentation/elbitar>
- [20] William Enck, Machigar Ongtang, and Patrick McDaniel. 2009. On Lightweight Mobile Phone Application Certification. In *Proceedings of the 16th ACM Conference on Computer and Communications Security* (Chicago, Illinois, USA) (*CCS '09*). Association for Computing Machinery, New York, NY, USA, 235–245. <https://doi.org/10.1145/1653662.1653691>
- [21] Zheran Fang, Weili Han, and Yingjiu Li. 2014. Permission based Android security: Issues and countermeasures. *computers & security* 43 (2014), 205–218.
- [22] Johannes Feichtner and Stefan Gruber. 2020. Understanding Privacy Awareness in Android App Descriptions Using Deep Learning. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy* (New Orleans, LA, USA) (*CODASPY '20*). Association for Computing Machinery, New York, NY, USA, 203–214. <https://doi.org/10.1145/3374664.3375730>
- [23] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*. 1–14.
- [24] Thomas Franke, Christiane Attig, and Daniel Wessel. 2019. A personal resource for technology interaction: development and validation of the affinity for technology interaction (ATI) scale. *International Journal of Human-Computer Interaction* 35, 6 (2019), 456–467.
- [25] Frederik Funke and Ulf-Dietrich Reips. 2012. Why semantic differentials in web-based research should be made from visual analogue scales and not from 5-point scales. *Field methods* 24, 3 (2012), 310–327.
- [26] Marco Furini, Silvia Mirri, Manuela Montangelo, and Catia Prandi. 2020. Privacy perception when using smartphone applications. *Mobile Networks and Applications* 25, 3 (2020), 1055–1061. <https://doi.org/10.1007/s11036-020-01529-z>
- [27] Hongcan Gao, Chenkai Guo, Yanfeng Wu, Naipeng Dong, Xiaolei Hou, Sihan Xu, and Jing Xu. 2019. AutoPer: Automatic Recommender for Runtime-Permission in Android Applications. In *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, Vol. 1. IEEE, 107–116. <https://doi.org/10.1109/COMPSAC.2019.00024>
- [28] Gabriella M Harari. 2020. A process-oriented approach to respecting privacy in the context of mobile phone tracking. *Current opinion in psychology* 31 (2020), 141–147. <https://doi.org/10.1016/j.copsyc.2019.09.007>
- [29] Gunnar Harboe and Elaine M. Huang. 2015. Real-World Affinity Diagramming Practices: Bridging the Paper-Digital Gap. In *Proc. 33rd Annual ACM Conf. Human Factors in Computing Systems*. ACM, New York, NY, USA, 95–104. <https://doi.org/10.1145/2702123.2702561>

- [30] Jason I Hong, Yuvraj Agarwal, Matt Fredrikson, Mike Czapik, Shawn Hanna, Swarup Sahoo, Judy Chun, Won-Woo Chung, Aniruddh Iyer, Ally Liu, et al. 2021. The Design of the User Interfaces for Privacy Enhancements for Android. *arXiv preprint arXiv:2104.12032* (2021).
- [31] Jason I Hong, Yuvraj Agarwal, Matt Fredrikson, Mike Czapik, Shawn Hanna, Swarup Sahoo, Judy Chun, Won-Woo Chung, Aniruddh Iyer, Ally Liu, et al. 2021. *Designing Privacy-Enhanced Android's User Interfaces*. Technical Report. Carnegie Mellon University.
- [32] Jinseong Jeon, Kristopher K Micinski, Jeffrey A Vaughan, Ari Fogel, Nikhilesh Reddy, Jeffrey S Foster, and Todd Millstein. 2012. Dr. android and mr. hide: fine-grained permissions in android applications. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*. 3–14.
- [33] Aycan Kaya, Reha Ozturk, and Cigdem Altin Gumussoy. 2019. Usability measurement of mobile applications with system usability scale (SUS). In *Industrial Engineering in the Big Data Era: Selected Papers from the Global Joint Conference on Industrial Engineering and Its Application Areas, GJCIE 2018, June 21–22, 2018, Nevsehir, Turkey*. Springer, 389–400.
- [34] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. 2012. A conundrum of permissions: installing applications on an android smartphone. In *Financial Cryptography and Data Security: FC 2012 Workshops, USEC and WECSR 2012, Kralendijk, Bonaire, March 2, 2012, Revised Selected Papers 16*. Springer, 68–79.
- [35] Paul E Ketelaar and Mark Van Balen. 2018. The smartphone as your follower: The role of smartphone literacy in the relation between privacy concerns, attitude and behaviour towards phone-embedded tracking. *Computers in Human Behavior* 78 (2018), 174–182. <https://doi.org/10.1016/j.chb.2017.09.034>
- [36] Florian Keusch, Bella Struminskaya, Christopher Antoun, Mick P Couper, and Frauke Kreuter. 2019. Willingness to participate in passive mobile data collection. *Public opinion quarterly* 83, S1 (2019), 210–235. <https://doi.org/10.1093/poq/nfz007>
- [37] Florian Künzler. 2019. Context-aware notification management systems for just-in-time adaptive interventions. In *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 435–436. <https://doi.org/10.1109/PERCOMW.2019.8730874>
- [38] Nicholas D Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury, and Andrew T Campbell. 2010. A survey of mobile phone sensing. *IEEE Communications magazine* 48, 9 (2010), 140–150.
- [39] Hansoo Lee, Joonyoung Park, and Uichin Lee. 2022. A Systematic Survey on Android API Usage for Data-Driven Analytics with Smartphones. *ACM Comput. Surv.* 55, 5, Article 104 (dec 2022), 38 pages. <https://doi.org/10.1145/3530814>
- [40] Yuanchun Li, Yao Guo, and Xiangqun Chen. 2016. PERUIM: Understanding Mobile Application Privacy with Permission-UI Mapping. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (Heidelberg, Germany) (UbiComp '16)*. Association for Computing Machinery, New York, NY, USA, 682–693. <https://doi.org/10.1145/2971648.2971693>
- [41] Andrew Lowe, Anthony C Norris, A Jane Farris, and Duncan R Babbage. 2018. Quantifying thematic saturation in qualitative data analysis. *Field methods* 30, 3 (2018), 191–207.
- [42] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.
- [43] Vikas Kumar Malviya, Chee Wei Leow, Ashok Kasthuri, Yan Naing Tun, Lwin Khin Shar, and Lingxiao Jiang. 2023. Right to Know, Right to Refuse: Towards UI Perception-Based Automated Fine-Grained Permission Controls for Android Apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (Rochester, MI, USA) (ASE '22)*. Association for Computing Machinery, New York, NY, USA, Article 186, 6 pages. <https://doi.org/10.1145/3551349.3559556>
- [44] Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L. Mazurek, and Jeffrey S. Foster. 2017. User Interactions and Permission Use on Android. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 362–373. <https://doi.org/10.1145/3025453.3025706>
- [45] Mohammad Nauman, Sohail Khan, and Xinwen Zhang. 2010. Apex: extending android permission model and enforcement with user-defined runtime constraints. In *Proceedings of the 5th ACM symposium on information, computer and communications security*. 328–332.
- [46] Katarzyna Olejnik, Italo Dacosta, Joana Soares Machado, Kevin Huguénin, Mohammad Emtiyaz Khan, and Jean-Pierre Hubaux. 2017. Smarper: Context-aware and automatic runtime-permissions for mobile devices. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1058–1076. <https://doi.org/10.1109/SP.2017.25>
- [47] Emmanuel Onu, Michael Mireku Kwakye, and Ken Barker. 2020. Contextual privacy policy modeling in iot. In *2020 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech)*. IEEE, 94–102.
- [48] John O'donoghue and John Herbert. 2012. Data management within mHealth environments: Patient sensors, mobile devices, and databases. *Journal of Data and Information Quality (JDIQ)* 4, 1 (2012), 1–20.
- [49] Paul Pearce, Adrienne Porter Felt, Gabriel Nunez, and David Wagner. 2012. Addroid: Privilege separation for applications and advertisers in android. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security*. 71–72.
- [50] Jan Pennekamp, Martin Henze, and Klaus Wehrle. 2017. A survey on the evolution of privacy enforcement on smartphones and the road ahead. *Pervasive and Mobile Computing* 42 (2017), 58–76.
- [51] Sarah Prange, Sven Mayer, Maria-Lena Bittl, Mariam Hassib, and Florian Alt. 2021. Investigating User Perceptions Towards Wearable Mobile Electromyography. In *Human-Computer Interaction—INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part IV 18*. Springer, 339–360.
- [52] Yiting Qu, Suguo Du, Shaofeng Li, Yan Meng, Le Zhang, and Haojin Zhu. 2020. Automatic permission optimization framework for privacy enhancement of mobile applications. *IEEE Internet of Things Journal* 8, 9 (2020), 7394–7406. <https://doi.org/10.1109/JIOT.2020.3039472>

- [53] Ulf-Dietrich Reips and Frederik Funke. 2008. Interval-level measurement with visual analogue scales in Internet-based research: VAS Generator. *Behavior research methods* 40, 3 (2008), 699–704.
- [54] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-Scale Assessment of Mobile Notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 3055–3064. <https://doi.org/10.1145/2556288.2557189>
- [55] Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. 2017. Identifying the provision of choices in privacy policy text. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 2774–2779. <https://doi.org/10.18653/v1/D17-1294>
- [56] Albrecht Schmidt, Michael Beigl, and Hans Gellersen. 1999. There is more to context than location. *Computers & Graphics* 23, 6 (1999), 893–901. [https://doi.org/10.1016/S0097-8493\(99\)00120-X](https://doi.org/10.1016/S0097-8493(99)00120-X)
- [57] Gian Luca Scoccia, Ivano Malavolta, Marco Autili, Amleto Di Salle, and Paola Inverardi. 2019. Enhancing trustability of android applications via user-centric flexible permissions. *IEEE Transactions on Software Engineering* 47, 10 (2019), 2032–2051. <https://doi.org/10.1109/TSE.2019.2941936>
- [58] John S Seberger, Marissel Llavore, Nicholas Nye Wyant, Irina Shklovski, and Sameer Patil. 2021. Empowering resignation: There’s an app for that. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18. <https://doi.org/10.1145/3411764.3445293>
- [59] Bingyu Shen, Lili Wei, Chengcheng Xiang, Yudong Wu, Mingyao Shen, Yuanyuan Zhou, and Xinxin Jin. 2021. Can Systems Explain Permissions Better? Understanding Users’ Misperceptions under Smartphone Runtime Permission Model. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 751–768. <https://www.usenix.org/conference/usenixsecurity21/presentation/shen-bingyu>
- [60] Gulshan Shrivastava, Prabhat Kumar, Deepak Gupta, and Joel JPC Rodrigues. 2020. Privacy issues of android application permissions: A literature review. *Transactions on Emerging Telecommunications Technologies* 31, 12 (2020), e3773. <https://doi.org/10.1002/ett.3773>
- [61] Christoph Stach and Bernhard Mitschang. 2013. Privacy management for mobile platforms—a review of concepts and approaches. In *2013 IEEE 14th International Conference on Mobile Data Management*, Vol. 1. IEEE, 305–313.
- [62] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 168, 24 pages. <https://doi.org/10.1145/3544548.3581060>
- [63] Christopher Thompson, Maritza Johnson, Serge Egelman, David Wagner, and Jennifer King. 2013. When It’s Better to Ask Forgiveness than Get Permission: Attribution Mechanisms for Smartphone Resources. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (Newcastle, United Kingdom) (SOUPS '13). Association for Computing Machinery, New York, NY, USA, Article 1, 14 pages. <https://doi.org/10.1145/2501604.2501605>
- [64] Lynn Tsai, Primal Wijesekera, Joel Reardon, Irwin Reyes, Serge Egelman, David Wagner, Nathan Good, and Jung-Wei Chen. 2017. Turtle guard: Helping android users apply contextual privacy preferences. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. 145–162.
- [65] Na Wang, Bo Zhang, Bin Liu, and Hongxia Jin. 2015. Investigating effects of control and ads awareness on android users’ privacy behaviors and perceptions. In *Proceedings of the 17th international conference on human-computer interaction with mobile devices and services*. 373–382.
- [66] Yang Wang, Huichuan Xia, and Yun Huang. 2016. Examining American and Chinese internet users’ contextual privacy preferences of behavioral advertising. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 539–552.
- [67] Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. 2018. Dynamically regulating mobile application permissions. *IEEE Security & Privacy* 16, 1 (2018), 64–71.
- [68] Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman. 2018. Contextualizing Privacy Decisions for Better Prediction (and Protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173842>
- [69] Maximiliane Windl, Niels Henze, Albrecht Schmidt, and Sebastian S Feger. 2022. Automating Contextual Privacy Policies: Design and Evaluation of a Production Tool for Digital Consumer Privacy Awareness. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [70] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [71] Kuang-Wen Wu, Shaio Yan Huang, David C Yen, and Irina Popova. 2012. The effect of online privacy policy on consumer privacy concern and trust. *Computers in human behavior* 28, 3 (2012), 889–897. <https://doi.org/10.1016/j.chb.2011.12.008>
- [72] Yajin Zhou, Xinwen Zhang, Xuxian Jiang, and Vincent W Freeh. 2011. Taming information-stealing smartphone applications (on android). In *Trust and Trustworthy Computing: 4th International Conference, TRUST 2011, Pittsburgh, PA, USA, June 22-24, 2011. Proceedings 4*. Springer, 93–107.