

Article

Contrastive vs. Example-Based Explanations: Designing for Better User Comprehension in Smartphone Privacy Interfaces

Ananya Bhadauria ^{1,*}  and Andreas Riener ^{2,*} ¹ Independent Researcher, 68789 Sankt Leon-Rot, Germany² Human-Computer Interaction Group, Technische Hochschule Ingolstadt, Esplanade 10, 85049 Ingolstadt, Germany

* Correspondence: 4nanyabhadauria@gmail.com (A.B.); andreas.riener@thi.de (A.R.)

Abstract

Smartphone applications routinely collect and share personal data, yet users often struggle to understand these practices, particularly when third-party data sharing is involved. Existing mechanisms such as privacy policies and notices provide limited support for user understanding. To address this gap, we investigated how explainability concepts can enhance contextual privacy policies for mobile apps. We designed two interface prototypes integrating explanation strategies: contrastive explanations, which clarify data-sharing boundaries, and example-based explanations, which illustrate counterfactual scenarios. In an exploratory between-subjects user study ($N = 30$), we evaluated their impact on comprehensibility, simplicity, cognitive load, and experience. Statistical analysis revealed no significant differences between the two types. Hence, the findings should be read as descriptive trends rather than confirmatory effects. These trends suggested example-based explanations better supported users mental models, while contrastive explanations better supported decision-making. Our findings contribute design recommendations on applying explanation strategies in privacy interfaces, offering guidance for developers and researchers seeking to improve user understanding and trust in digital systems.

Keywords: smartphone privacy; explainability; contextual privacy policies; privacy interfaces

1. Introduction

Smartphones have become more than just a communication tool due to the integration of various applications that users install on their mobile phones [1,2]. However, user interaction with mobile applications inevitably involves the exchange of personal information [1,3]. Data protection laws such as the European General Data Protection Regulation (GDPR) [4] mandate companies to inform users what data is gathered about them, what is stored, and how and by whom these data are being used. Modern smartphone operating systems, such as Android and iOS, incorporate Runtime Privacy Notices (RPNs), enabling users to control app access to personal data such as photos and contacts [5]. However, due to several limitations, these notices fail to effectively explain to users data practices in a transparent manner, resulting in uninformed consent [6,7]. End-users consequently find this information difficult to understand, even though they have the right to clear and easily accessible explanations of how their data is collected, used, and shared [8]. The runtime permission model presents several challenges, starting with the use of brief permission descriptions. These short prompts make it difficult for users to grasp which types of personal data are being accessed [6].



Academic Editor: Richard Jiang

Received: 25 April 2026

Revised: 18 June 2026

Accepted: 20 June 2026

Published: 26 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

To address this limitation, the concept of Contextual Privacy Policies (CPPs) has been developed, which increases privacy awareness among users by providing context-relevant privacy information [9]. CPPs break down privacy policies into smaller, more digestible pieces, shown only within relevant contexts in mobile apps. To be relevant, contextual privacy notices should focus on unanticipated data practices, as the disclosure of personal information to third-party entities. [9,10]. Although CPPs are promising, creating accurate contextual notices for mobile applications presents unique challenges [9]. Companies such as Apple or Facebook have started implementing contextual privacy notices, but their effectiveness in improving user understanding of privacy information remains uncertain [10].

To further enhance transparency in digital systems, it is crucial to develop innovative approaches that provide users with clear and comprehensible privacy information. This need aligns with the concept of explainability, which focuses on improving the user's understanding of how a system operates [11,12]. Explanations can help users understand the system's behavior, the reasoning behind certain decisions, and how internal processes work. This can improve users' understanding of privacy aspects and support more informed decisions [8]. While privacy and explainability aspects have been well studied individually, their integration remains underexplored. This gap is especially critical in domains where both are essential and must be addressed together [13].

To address this gap, we began our research to find how we can design explainable privacy interfaces that enhance user understanding of app-specific data practices. To answer this, first, we examined the current level of user understanding regarding application-specific data practices, such as data collection, data usage, and third-party sharing, through a pre-study survey. This survey provided insights into user comprehension, highlighting that users exhibit incomplete and uneven comprehension, with third-party data sharing representing as the most significant knowledge gap.

Second, we investigate how users perceive current privacy explanations in the literature when they are integrated into smartphone applications. Our findings reveal that users perceive privacy explanations as effective for improving understanding of app-specific data practices. However, participants still perceived explanations of third-party data sharing as the least understood and emphasized the need for specific details on how their data is shared. They suggested presenting this information in a short, contextual manner that avoids overwhelming users, and combining visuals with text to make explanations quicker and easier to grasp.

Finally, building on these insights, we designed and evaluated two types of explainable privacy interfaces inspired from Explainable Artificial Intelligence (XAI) literature and integrated them into a smartphone application: one based on example-based explanation strategy and the other based on contrastive explanations. A comparative user study was conducted based on an explainability quality model [8] to assess their impact on user understanding. Although we found no statistically significant differences between the two explanation types across comprehension, simplicity, cognitive load, and user experience, descriptive trends offered useful insights. Our work is particularly relevant to the design of explainable privacy interfaces in smartphone applications. While both example-based and contrastive explanations effectively supported user understanding of data sharing practices, descriptive trends suggested that example-based explanations were associated with somewhat higher mental model accuracy, whereas contrastive explanations appeared more supportive of decision-making. Thus, it is crucial to design privacy interfaces that apply layered explanations, progressive disclosure, and transparent third-party communication while managing cognitive load through visual support.

Contribution Statement

Our contribution to the HCI community ([14], ch. 1.2) is

1. Artifact or system and method: We developed interactive prototypes of two explainable privacy interfaces to systematically evaluate the effectiveness of different explanation strategies.
2. Empirical study that tells us about how people use a system: We conducted an exploratory between-subjects study ($N = 30$) to examine the influence of the two types of explanations on user understanding of data sharing practices. Although not statistically significant, our results indicate a trend in which example-based explanations more effectively support users' mental models, while contrastive explanations more effectively support decision-making.

2. Background and Related Work

2.1. Privacy in Smartphone Contexts

Privacy, in the context of smartphone applications, has evolved from the classical notion of the right to be let alone [15] to a more nuanced understanding centered on user control and context-specific information flows [3,16]. Westin's work frames privacy as the individual's ability to decide when, how, and to what extent information is shared [3]. These definitions are particularly relevant in mobile ecosystems, where personal data is continuously collected and processed, often beyond the immediate awareness of users [15]. Drawing from these perspectives, privacy in the context of smartphone apps can be defined as the condition in which personal information generated or accessed by mobile applications flows in a manner consistent with user-understood, context-appropriate norms.

2.1.1. Existing Privacy Concerns in Smartphone Applications

Smartphone applications have become deeply embedded in daily life, but this integration has raised significant privacy concerns, as applications routinely collect and transmit sensitive personal data, often beyond what users expect or understand [17–19]. Users often lack awareness of what data is collected, used, and shared, as privacy notices and permission dialogues are frequently too complex or vague to support informed decisions, leading to a situation where most users cannot accurately recall which permissions they have granted or how to adjust them after installation [20]. In summary, privacy concerns in smartphone applications are multifaceted, involving excessive data collection, lack of transparency about third-party sharing, and insufficient user awareness.

2.1.2. Third-Party Sharing

Brunotte et al. [21] define third-party sharing as a privacy practice that involves either sharing data with external organizations or data collection directly by such entities. A third party is understood as any company or organization other than the first-party provider that owns the website or application. Nguyen et al. [22] conducted a comprehensive analysis of over 239,000 Android apps and found that a substantial portion transmitted users' personal data to third-party data controllers, often for advertising purposes, without implementing any form of user consent notice. Schaub et al. [23] demonstrated that users often fail to comprehend descriptions of third-party services in app privacy notices unless they already recognize the company names. This approach falls short of the notice's purpose of teaching users about data practices, because simply naming a company or using a vague label like "data reseller" or "government agency" doesn't explain what happens to their personal information. Third-party data sharing is now widespread in mobile apps, but privacy notices usually only list the names or general types of these partners. As a result, users don't really understand how or why their personal information is being used [3].

2.2. Privacy Notices in Smartphone Applications

Privacy notices are the primary means by which information about privacy policies is communicated to users. These notices take on various forms, ranging from general privacy policy documents to permission dialogues in mobile applications [7]. In this section, we discuss the different types of privacy notices that are available to end-users in smartphone applications that inform users about data practices.

2.2.1. Privacy Policies

Privacy policies are a standard method for communicating data practices to users, outlining what information is collected, how it is used, how information is shared, and what choices users have to manage these activities [24]. However, Windl et al. [24] highlight that privacy policies are often lengthy, written in complex legal language, making them inaccessible and difficult for users to read or understand. Ebert et al. [10] reinforce this finding, noting that privacy policies primarily serve legal compliance and liability limitation for companies, rather than genuinely informing users. McDonald and Cranor [25] calculated that, over the course of a year, the average user would need to spend approximately 201 h reading all the privacy policies they encounter. This striking figure highlights the significant time burden imposed by current privacy policy practices, illustrating why most users are unable to fully review and understand the terms presented to them. Studies show that most users do not read privacy policies at all, and those who do spend an average of only 73 s per policy which is far too little time to grasp the content [24]. This lack of engagement with privacy policies increases the risk of misunderstandings regarding the provider's data practices [10]. In summary, substantial improvements are required, particularly regarding the length, wording, and visibility of privacy policies to create policies that are truly accessible and support meaningful notice and consent for users [9,10,24].

2.2.2. Runtime Privacy Notices

Smartphone operating systems Android and iOS, introduced runtime privacy notices as a response to the limitations of traditional static privacy policies, which are often lengthy and difficult for users to understand [9]. Run-time Privacy (RPNs) are used to enable users to make informed choices about granting or denying application requests for access to personal data such as photos or contact information [6].

These notices typically appear as permission prompts at the moment when sensitive data is accessed, providing users with timely and context-specific information through a just-in-time disclosure mechanism [9]. Although runtime permissions appear when an app first requests access to a sensor and are thus contextual in a temporal sense, they offer only a momentary snapshot of data access. Prior work shows that these notices surface only at install time or during first invocation, giving no insight into how data may later be transmitted, shared, stored, or repurposed [9]. Pan et al. [9] identify several key limitations that restrict their effectiveness in educating users about data sharing. First, while the prompts appear at the moment when permissions are granted, they only provide a narrow snapshot of the immediate access request (e.g., camera or location) and omit any information about how collected data may subsequently be transmitted off-device, shared with third parties, stored, or repurposed. Users thus gain no insight into the broader life cycle of their data beyond the initial ask. From such short descriptions, users cannot completely comprehend the risks of allowing these permissions [6]. Shen et al. [6] also confirm that these runtime notices often fall short in addressing user's specific concerns, such as identifying the third parties involved in data sharing or the frequency of data access. Moreover, existing implementations of just-in-time permission prompts are often ineffective, as security-related warnings are typically detached from the context in which

data access occurs. This disconnection reduces their ability to capture users' attention and foster meaningful awareness.

2.2.3. Contextual Privacy Policies

Contextual privacy policies (CPPs) have emerged as a promising approach to address the longstanding shortcomings of traditional privacy notices. By embedding concise, context-relevant information directly within the user interface, CPPs aim to improve user awareness and understanding of data practices at the moment they matter most. This concept introduces a novel approach by breaking down traditional privacy notices into concise, easy-to-understand segments that are presented only within relevant usage contexts in mobile applications [9]. These privacy-related contexts are tied to specific elements of the Graphical User Interface (GUIs) where data interactions occur [9]. Most studies [9,10,24,26], have revealed the benefits of CPPs in making privacy information understandable to users. By presenting privacy information in-situ, such as alongside a login form or location request. CPPs reduce cognitive load and make the relevance of data practices more apparent to users [26]. Human evaluations report that users rate CPPs as more accessible and are more willing to engage with them compared to traditional policies [9,24].

Windl et al. [24] note that users rate CPPs as more accessible and less overwhelming than traditional privacy policies. The authors highlights that while CPPs show promise in enhancing user awareness and decision-making, several key improvements are needed for their content and presentation to be truly effective. Their qualitative study found that CPPs increased participant's privacy awareness, often prompting more reflective behavior and, in some cases, immediate actions such as leaving a website when unexpected data sharing was revealed. However, users identified persistent barriers, including information overload, lack of time, and the complexity of even contextualized snippets, which sometimes remained too lengthy or abstract for comprehension [24]. In conclusion only providing privacy policies in relevant context alone is not enough to achieve complete understanding of data practices for end-users, we need to go beyond just presenting information to users and also focus on making the content understandable. To bridge this gap, there is a need to identify strategies which can improve user understanding of data practices, while being relevant to the concept of Contextual Privacy Policies.

2.3. Explainability to Improve Understandability

As discussed above (Section 2.2), privacy notices in smartphone apps face numerous challenges that limit their effectiveness. They often fail to capture users' attention, provide sufficient information, or facilitate understanding of data practices. As a result, users require a new approach to achieve clarity and transparency regarding their privacy [16]. There is a growing need for innovative methods that convey privacy information in a clear and understandable manner. This demand aligns closely with the principles of explainability, which aim to enhance user comprehension of system behavior [12]. In this section, we further investigate how explainability can improve these privacy notices in terms of enhancing user understanding of data practices.

In simple terms, explainability refers to a software system's capability to clarify its own behavior. Within requirements engineering, it is considered a non-functional requirement (NFR) of the system [3,11,27]. Explainability is an (NFR) that is essential for achieving broader software qualities such as Transparency [11,27,28]. By incorporating explainability, system opacity can be reduced, enabling users to understand why a particular output was generated and making it easier for them to make more informed decisions. Explanations are offered with the intent that the recipient comprehends how the software functions [28].

Therefore, explainability is also closely connected to understandability [3]. Therefore, an explainable system must be understandable, which also implies that it must be transparent in nature. Chazette and Schneider [29] highlight that explainability, while generally beneficial, can also actually reduce system understandability if the explanations are too complex, overly detailed, or not aligned with the user's background and needs. In summary, transparency is essential for delivering effective explanations to end-users. However, it must be tailored appropriately to the target audience to ensure that the explanations are understandable and help users better grasp how the system works. This also highlights the need to integrate transparent explanations to enhance user understanding of data practices.

2.3.1. Privacy Explanations as a Solution

In the context of privacy, explainability can be applied through explanations that help users understand how their personal data will or will not be processed by the system [16]. Further, we describe how privacy explanations can be a solution to address users lack of understanding about data practices. Privacy explanation serves as the medium through which privacy-related information is communicated to the end-user. As a result, privacy explanations that can clarify why an application requires access to a certain feature may positively influence how end users perceive the system [12].

2.3.2. Related Work

Explainability has been a central concern in the field of Artificial Intelligence (AI), particularly within the research area of Explainable Artificial Intelligence (XAI) [16]. Chazette et al. [11] conducted a study on explainability within the context of a navigation application and found that providing explanations can enhance usability. The authors also suggest users prefer to have the flexibility to request explanations when needed, pointing to a strong expectation for user autonomy and context-sensitive delivery. However, their findings also indicate that including excessive or irrelevant information may negatively impact the user experience, such as cognitive overload, visual clutter, or users feeling a loss of control. Brunotte et al. [16] study focused on examining how privacy explanations might contribute to building end-user trust in a system and enhancing user awareness of data privacy practices. In the study participants expressed discomfort regarding excessive collection of data, such as software tracking various aspects of their daily lives. Greatest discomfort and concern arose from excessive or unclear data collection, invasive advertising, uncertainty about third-party data sharing, and lack of transparency over data storage and use. According to the study conducted by Brunotte et al. [16] only 8.4% ($N = 155$) users actually read a privacy policy, this can also be a reason users feel discomfort and concern about their data collection and sharing. This finding highlights the importance of making privacy policies and notices accessible and understandable for end-users. The findings highlight that privacy explanations can enhance user trust and awareness by increasing transparency, helping users make informed decisions, and improving their sense of security when using a system.

2.4. Types of Explanations

Contrastive and example-based explanations are two influential approaches in the literature on explainable artificial intelligence XAI, each offering different strategies for making system decisions more understandable to users [3,30,31]. These two special types of explanations are identified that deviate from regular textual explanations [3].

2.4.1. Contrastive Explanations

In the field of XAI, contrastive explanations are introduced as an alternative to standard explanations. A regular explanation [16] clarifies why an event occurred, whereas a

contrastive explanation highlights why that event happened rather than another possible outcome. Miller [32] states that the complexity of intelligent software systems can make full regular explanations overwhelming for users. Providing a point of comparison through contrastive explanations helps reduce this cognitive load. This method reflects how humans typically reason, as they tend to question why one specific event happened instead of another anticipated or alternative outcome. Studies show that contrastive explanations enable users to pinpoint which factors influence specific outcomes, thereby enhancing their understanding of how the system makes individual decisions [30].

2.4.2. Example-Based Explanations

Another way to communicate information to users is through example-based explanations. In the field of XAI, these have emerged as an effective alternative to traditional textual explanations [3]. In automated decision systems, these explanations illustrate what changes would need to occur in order to alter the system's decision. In other words, they help users understand how a different outcome could be achieved by modifying certain inputs or conditions [33].

2.5. Knowledge Gap and Research Questions

As discussed above, prior work on smartphone privacy mechanisms, such as policies, runtime dialogues, and contextual notices, offers only partial improvements, as users still struggle to understand how their data is collected, shared, and processed. These shortcomings highlight the need for approaches that extend beyond mere compliance and actively promote understanding. Explainability, widely studied in AI systems, addresses similar challenges of transparency and user understanding, yet its potential in privacy interfaces remains underexplored.

This gap motivates our central research question (RQ): "How can explainable privacy interfaces be designed to enhance user understanding of app-specific data practices?", and is explored through three sub-questions:

- Sub—RQ1: What is the current level of user understanding and interest regarding app-specific data practices, such as data collection, usage, and third-party sharing?
- Sub—RQ2: How do users perceive current explanation with respect to understanding of app-specific data practices, and what improvements do they suggest?
- Sub—RQ3: How do different explainable privacy interfaces (example-based vs. contrastive) influence user understanding of app-specific data practices?

3. Methodology

To investigate how explainable privacy interfaces can enhance user understanding of app-specific data practices, we adopted a mixed-methods approach embedded in a User-Centered Design (UCD) process [34] and structured our research into four phases:

- Phase 1—Understanding the Problem Space: We started with an extensive literature review to identify limitations of existing privacy notices (see Section 2). These insights helped formulate initial research questions and inform later empirical designs.
- Phase 2—Eliciting User Needs: In this phase, we conducted a pre-study to evaluate how individuals perceive current privacy interfaces (see Section 4). A mixed-method survey was used to uncover gaps in comprehension and gather user feedback on explanation styles. The goal was to explore user preferences and identify specific pain points in understanding app data practices.
- Phase 3—Designing Explainable Privacy Interfaces: Drawing on insights from earlier phases, we developed two interface prototypes, one using contrastive explanations

and the other using example-based explanations (see Section 5). The designs were implemented as interactive prototypes for testing.

- Phase 4—Evaluation through Mixed-Methods Testing: In the final phase, we conducted a comparative user study to evaluate the impact of the two interface types on user understanding. Both quantitative data (e.g., comprehension scores, cognitive load) and qualitative insights (e.g., user feedback) were analyzed to inform discussion (see Section 6) and conclusions (see Section 7).

4. Survey Design

We ensured that the survey structure was directly aligned with its related sub-research questions (sub-RQ1 and sub-RQ2), ensuring that each section mapped to specific constructs of user understanding, interest, and perception regarding privacy explanations. It started with information on how participant data would be used and consisted of four sections with a total of 23 questions (see Appendix A). The structure of the survey, along with the sections addressing the respective research questions, is illustrated in Figure 1. In the pre-study, we operationalized “understanding” as participants’ perceived understanding of app-specific data practices. Perceived understanding reflects users’ self-reported familiarity and confidence with how apps collect, use, and share personal data. This approach follows prior privacy-explanation survey research, such as Brunotte et al. [16], which relies on perception-based measures to identify users’ awareness gaps before and after exposure to explanations.




Figure 1. Overview of the survey structure.

We conducted a pre-test with five participants to assess the survey’s quality. Two participants were PhD students from our university, and three candidates were master’s students from a non-technical department. Based on the pilot test, we made some corrections (e.g., the wording of the questions was improved to facilitate better understanding, and the order of questions was changed for improved flow).

4.1. Data Collection and Analysis

Data collection took place over a period of one month, starting in February 2025. Responses were collected via a web-based questionnaire created in LimeSurvey and hosted on the university server. The survey was distributed through multiple channels, including university mailing lists, LinkedIn, WhatsApp groups, and the personal network of the authors. In accordance with our university’s policy, no IRB approval was required for this study; however, it was conducted in full compliance with established ethical standards

(Declaration of Helsinki [35]): Participation was voluntary, participants could withdraw at any time, and provide comments or feedback. No physical or psychological harm occurred during the study. Inclusion criteria required owning a smartphone and having sufficient English proficiency. A quantitative analysis was conducted for Sections 1–3 using single-choice, multiple-choice, and 5-point Likert-scale items. The resulting data was organized and analyzed in Microsoft Excel.

We also analyzed open-ended responses by identifying recurring themes and calculating their frequency. These themes were first visualized using treemap diagrams to illustrate their relative prominence. To prioritize design focus, we then mapped the themes onto a 2×2 matrix with “Impact” on the Y-axis and “Frequency” on the X-axis [36].

4.2. Results—Pre-Study Survey

4.2.1. Demographics

We received a total of 154 responses. However, out of the 154 responses, 54 respondents did not complete the full questionnaire and provided answers only to the initial demographic section. These partial responses were excluded, resulting in 100 valid and complete responses. Gender distribution was fairly balanced, comprising approximately 45% female and 54% male respondents, with one participant (1%) not specifying gender. A significant portion of the participants were students (68%). This indicates that the participant pool was predominantly composed of students, suggesting that the findings may largely reflect the perceptions and experiences of an academically engaged population.

4.2.2. General Data About Smartphone Usage

All participants in the final sample were smartphone owners, out of which 42% of respondents reported using Android devices, while 58% used iOS. Most participants (95%) reported regularly using social media applications, followed closely by banking and finance apps (66%) and streaming or entertainment apps (65%). The widespread adoption of these app categories highlights their central role in participants’ daily digital activities, underscoring the relevance of privacy and data handling practices within these commonly used domains.

4.2.3. Sub—RQ1: Understanding and Interest in Data Practices

To assess participants’ current understanding of app-specific data practices, they were asked to self-report their understanding of these practices. Results indicate that the users have a limited general understanding of how applications handle data, particularly in relation to data sharing. As shown in Figure 2, 44% of the participants selected “Do not understand at all” or “Slightly understand” for data sharing. Low self-reported understanding was also observed for data usage (30%) and data collection (20%). Confidence in understanding varies across categories, with 30% of respondents feeling “Mostly” or “Fully understand” with regard to data sharing, 40% for data usage, and 51% for data collection. This suggests that while some familiarity exists, user understanding is often incomplete or uncertain.

To analyze whether participants are interested in learning more about how their data is collected, used, and shared, we asked participants about their interest in learning more. In total, 76% indicated they are “Mostly” or “Fully interested” in further information about data sharing, 69% for data usage, and 67% for data collection. We included the self-reported understanding and interest measures to identify which data practices users understand least and in which they are most interested, thereby informing the focus of our design. The results show uneven comprehension, with data sharing emerging as the least understood practice. This is particularly relevant for our study, as the main user study investigates explanations for data sharing. This pattern aligns with prior work

showing that users struggle to understand third-party data flows [16]. The interest ratings complement these findings. This combination of low perceived understanding and high interest supports the need for clearer, more accessible explanations and validates our decision to focus the main study on data-sharing practices.




Figure 2. Participant perceived understanding of how apps collect, use, and share data. (The chart displays percentage distributions across Likert scale levels (1 = Do not understand at all, 5 = Fully understand), with a focus on data sharing to highlight the lowest perceived comprehension).

4.2.4. Sub—RQ2: Perceived Understanding of Current Privacy Explanation

Having established the baseline understanding and interest levels in Sub-RQ1, we now turn to Sub-RQ2 to examine whether users perceive privacy explanations as effective tools for addressing the knowledge gaps identified, what specific characteristics make these explanations most helpful in supporting user understanding of data practices, and finally, what improvements do users suggest to improve user understanding of data practices. The privacy explanation scenario was modified and adapted from the work of Brunotte et al. [16].

Survey participants were presented with the following hypothetical scenario to assess their perceived understanding: “You have downloaded a new smartphone app called ‘City-Explorer’ that helps users discover hidden attractions, local restaurants, and personalized sightseeing routes. When you open ‘CityExplorer’ for the first time, it requests access to your Location without providing any explanation.”

In the first step, we asked the participants how important it is for them that the app tells them why it needs to use their location. A majority of users (63%) indicated that explanations are either “Very important” or “Fully important,” with only a small fraction (3%) rating them as “Not at all important” and 11% viewing them as “Slightly important.” This underlines strong user demand for clear privacy explanations within app interactions. This validates the finding of Brunotte et al. [16], which indicates that users are highly interested in receiving privacy explanations.

In the second step, we provided a regular explanation regarding the use of their location data. In line with prior work showing that users often struggle to understand references to third-party services unless they already recognize the company names [23], we intentionally used a deliberately vague label (“advertisers”) in the explanation. This design choice allowed us to examine whether participants could infer the implications of third-party data sharing even when specific entities were not named, reflecting realistic conditions in many current privacy notices.

Explanation: “Now, imagine the app provides the following explanations”:

- Implications of allowing access: Enhances user experience by providing you with personalized recommendations.
Note: Location data may be shared with advertisers.
- Implications of denying access: You get generic recommendations.
Note: Location data will not be shared with advertisers.

We then asked participants how helpful the provided explanations were for understanding the implications of allowing access, denying access, and third-party sharing (Figure 3). For allowing access, 64% rated the explanation “Very” or “Fully helpful.” For denying access, 58% did so, and only 12% rated it “Not helpful at all” or “Slightly helpful.” In contrast, for third-party sharing, only 44% found the explanation “Very” or “Fully helpful,” while 37% rated it “Not helpful at all” or “Slightly helpful,” marking third-party sharing as the clearest pain point. This finding reinforces that without detailed descriptions about the third-party and sharing, participants do not find the privacy explanations very helpful for their understanding.




Figure 3. Perceived helpfulness of privacy explanations. (The chart displays percentage distributions across Likert scale levels (1 = Not helpful at all, 5 = Fully helpful), with a focus on third-party sharing to highlight the lowest perceived helpfulness).

However, we acknowledge that the location-based scenario used in this survey may appear obvious, as recommending nearby places is a common and expected reason for location access. This simplicity, while beneficial for ensuring that all participants could easily follow the scenario, also introduces limitations. However, this design choice was intentional. Our Sub-RQ2 aimed to probe users’ perceptions of privacy explanations in a setting that reflects a realistic and frequently encountered permission request. Moreover, results from Sub-RQ1 showed that location sharing practices were the most poorly understood and most desired area for further explanation among participants, making location sharing a relevant and justified focus for examining how explanations shape perceived understanding. We discuss implications of this scenario choice further in (Section 6.5).

Improvements Identified in the Explanation

Lastly, we asked participants to suggest improvements that would enhance understanding of data practices. To further explore the effect of explanation on user understanding of data practices, we asked participants open-text questions focusing on how the explanation affected their understanding or perception of the app’s use of their data.

To analyze the open-text responses, we conducted an inductive thematic analysis following Braun and Clarke’s [37] flexible framework. The analysis proceeded through several iterative steps. First, all responses were read repeatedly to achieve familiarization and build an initial sense of participant concerns. We then generated initial codes directly from

the data by marking segments of text that reflected meaningful observations or suggestions. These codes were inductively grouped into preliminary themes by clustering semantically similar ideas, and the emerging themes were refined by reviewing them against the full dataset to ensure that each theme accurately captured the underlying comments. Braun and Clarke [37] note that frequency counts can be used descriptively; therefore, once the themes were finalized, we counted the number of participants who mentioned each theme to indicate its relative prominence in the dataset. We report frequency counts only to indicate how commonly a theme appeared across participant responses, not as statistical evidence. All coding and theme development were conducted by a single researcher. This procedure provided a systematic yet flexible approach for identifying user-centered patterns in the survey responses.

Thematic analysis surfaced five priorities: (1) greater transparency about what, why, and how data is shared ($N = 43$), a participant stated: “List what exact location data is shared and with which advertisers.”; (2) visual aids to simplify complex flows ($N = 40$), another participant stated: “Use step-by-step visuals or infographics instead of all text at once.”; (3) shorter, contextual copy with progressive disclosure ($N = 16$); (4) legal/data assurances ($N = 6$); and (5) limited personalization options ($N = 2$). Overall, while our explanation, adapted from Brunotte et al. [16], helped, participants indicated it needs clearer third-party specifics, stronger visual support, and tighter, contextual presentation.

4.3. Summary and Implications for Design

Taken together, the pre-study survey provided two main insights. First, participants showed uneven understanding of app data practices, with third-party sharing emerging as the primary gap. At the same time, users expressed strong interest in receiving clearer explanations about how their data is collected, used, and shared.

Second, while privacy explanations improved perceived understanding to some extent, they did not fully resolve uncertainty, particularly for complex practices like third-party sharing. Based on these findings, we derived design guidelines (Table 1), prioritized via a 2×2 matrix of impact and frequency, which directly informed the prototype development in Section 5.

Table 1. Design guidelines derived from the pre-study.

Design Guideline	Key Inference from Pre-Study	Prioritization (Frequency/Impact)	Design Implication
Enhance transparency	Persistent uncertainty about data details, especially data sharing	High/High	Include “what, why, and how” breakdowns in explanations, the 2W1H principle [16]
Incorporate visual aids	Need for simplified presentation to clarify complexity	High/High	Integrate infographics or icons alongside text for better accessibility
Short and contextual explanations	Demand for concise, timely information without overload	Moderate/High	Deliver just-in-time notices during permission requests or app setup
Legal assurance and customization	Calls for compliance information and personalization	Low/Low	Add optional layers for detailed assurances or user-configurable views

5. Design of Explainable Privacy Interfaces

The following sections describe our design approach for the prototype development.

5.1. Dimensions of Privacy Notices

Using Schaub et al.’s [23] design space for privacy notices, we outlined our design choices along its four key dimensions. Its core components include timing (when the notice is shown), channel (the method of delivery), modality (the forms of interaction used),

and control (the way user choices are offered). These decisions formed the foundation for implementing the explanations in our prototypes.

- **Timing:** Notices are most effective when aligned with user attention [23]. Our design adopts a layered persistent notice that remains available while the data practice is active. This ensures contextual relevance and visibility without interrupting the flow of interaction.
- **Channel:** We deliver notices through the primary channel (on-screen), keeping them immediate and contextually relevant. This choice reflects pre-study findings that emphasized transparency and just-in-time presentation.
- **Modality:** A visual presentation with icons and concise text was selected, as it balances detail and simplicity, supporting user comprehension without introducing additional learning curves. This aligns with pre-study preferences for visual aids.
- **Control:** Instead of blocking mechanisms (as in iOS/Android RPNs), we implemented a non-blocking explanatory layer. This complements the existing consent flows of the current privacy notices, allowing users to explore details at their own pace.

5.2. Prototype Development

Given the pre-study's revelation that regular textual explanations, while somewhat helpful for basic implications, leave users with ongoing uncertainty about data practices (especially data sharing), there is a compelling need to explore verified methods from the literature to bolster understanding. To address this, we turned to XAI literature, which offers promising strategies for making complex systems more interpretable. As discussed in Section 2.4, contrastive and example-based explanations have been shown to aid user comprehension in data-related contexts by providing comparative insights and relatable instances [30,31], potentially increasing transparency when integrated into privacy interfaces. The following sections detail how we integrated these explanation types in our interface screens.

We embedded the privacy interface into Instagram (Android version, June 2025), a widely used social media app, ensuring familiarity for participants [38,39]. The content of our prototype screens was guided by a review of the official Instagram privacy policy (<https://privacycenter.instagram.com/>, last access: 25 April 2026). We particularly focused on the section "How do we share information with third parties?" of the policy, as our pre-study results highlight that data-sharing practices are the least understood by participants. We created two versions for our prototype. Both conditions used the same prototype interface and interaction flow. The only difference between versions was the explanation screen, which presented either the contrastive or the example-based explanation, while all other screens and UI elements remained identical.

As shown in Figure 4 (left image), Interface A is a contrastive explanation, which follows the principle that explanation are more understandable when they highlight not only what is shared, but also what is not shared by providing a direct point of comparison [3,30,31]. Thus, the interface clearly distinguishes between items like "Your approximate city-level location" (what is shared), and "Your precise GPS location" (what is not shared). Which, directly reference the data categories specified in the Instagram's privacy policy section about data-sharing practices.

Alternatively, as shown in Figure 4 (right image), Interface B is an example-based explanation, which presents outcomes based on two specific user actions, labeled as "If Allowed" and "If Denied." This aligns with the principle of example-based explanation from XAI, by providing clear, concrete scenarios [3,30,31]. These scenarios illustrate what would happen if a user chooses to share or not share specific data. For instance, the interface presents outcomes such as: "If you share your location, personalized ads will be shown

based on your nearby stores” vs. “If you don’t share, you will still see ads, but they won’t be location-based.” This approach focuses on outcomes rather than data items, aligning with the example-based principle of conveying meaning through situational examples.

In both explanation versions, the icons used to denote “shared” and “not shared” (or “allowed” and “denied”) were intentionally designed as status indicators rather than representational icons depicting data types or stakeholders. This choice ensured visual consistency across the two formats and reduced the interpretive burden that more complex icons might introduce. Our aim was to emphasize the distinction between disclosure states. What data leaves the device versus what remains internal, without requiring users to infer additional symbolic meaning in the explanations.

For completeness and to improve clarity regarding the interaction flow, all prototype screens used in the study are provided in Appendix C.




Figure 4. Comparison of interface A (contrastive explanation; left) and interface B (example-based explanation; right) for advertisers.

To evaluate our designs, we conducted a between-subject study. The goal was to investigate the impact of the two types of explanations on users’ understanding of data sharing practices. The following sections outline the study setup, measures, and results.

5.3. Quality Model

To effectively evaluate the explainability of our contrastive and example-based privacy interfaces, we drew upon established frameworks from the literature. In particular, we examined the research by Deters et al. [8]. The authors highlight that to define requirements for explainable systems, it is necessary to first understand the elements that contribute to a system’s explainability. This includes considering various aspects, such as understandability and transparency. In the following section, we describe the quality model used in

our study to formulate hypotheses and test the interfaces based on the criteria and aspects defined by the model.

5.4. Understandability

According to the Deters et al. [8] quality model, the first aspect of explainability is understandability. The major aim of understandability is that the addressee can easily understand explanations. Adapting the above quality model, we formulated our hypotheses by focusing on significant criteria under the understandability aspect, selecting measures that map to those criteria for comparing the Interface A (contrastive) and Interface B (example-based) interfaces. As shown in Figure 5, we evaluated the explanations across three main criteria: (1) comprehensibility of presentation form, (2) cognitive load, and (3) simplicity, using the established metrics outlined in the model. Comprehension was assessed through mental model accuracy. We define mental model accuracy as the extent to which a user's internal understanding of how a system operates corresponds to the system's actual functioning. Prior work shows that mental models can be assessed by asking users to explain how they believe a system works and comparing these to the actual system behavior. Additionally, subjective understandability reflects how understandable a data practice feels to users [8]. Cognitive load was measured with the NASA-TLX [40]. Simplicity was examined using the identification of irrelevant items. In addition, we employed the UEQ-S [41] to capture overall user experience.




Figure 5. Adapted quality model for measuring understandability aspect of explainability in our privacy explanations [8].

Following are the formulated Hypotheses (H1–H4) based on the quality model to answer the Sub—RQ3: “How do different explainable privacy interfaces (example-based vs. contrastive) influence user understanding of app-specific data practices?”

- H_1 : There is a significant difference in user comprehension between Interface A and Interface B, as measured by mental model accuracy.
- H_2 : There is a significant difference in perceived simplicity between Interface A and Interface B, as measured by the identification of irrelevant elements in the explanation.
- H_3 : There is a significant difference in perceived cognitive load between Interface A and Interface B, as measured by the NASA-TLX scale.
- H_4 : There is a significant difference in user experience between Interface A and Interface B, as measured by the User Experience Questionnaire (UEQ-S).

5.5. Participant Criteria and Setting

Participants were required to own a smartphone and be familiar with Instagram application. We recruited a total of 30 participants (15 per condition) for our study. A formal a-priori power analysis was not conducted before data collection. The sample size ($N = 30$) was chosen in line with common practice in HCI and UX research for between-subject designs. The sample size limits statistical power for detecting small effects, but is appropriate for identifying large usability differences [42]. We position this work as an exploratory HCI/UX study rather than a confirmatory experiment. Small-sample exploratory studies of this kind are common and accepted in HCI and UX research, where in-depth sessions with a limited number of participants are routinely used to evaluate early-stage interface concepts [43]. We therefore report our quantitative results as descriptive trends and use them to motivate larger, adequately powered studies. The study was conducted remotely via Microsoft Teams, participants were given prototype control through AnyDesk (June 2025 version). Sessions were moderated and participants were guided and provided with the tasks one after the other [44]. Recruitment was done via social media and the university mailing list, and sessions were scheduled with Calendly (June 2025 version). Names and emails collected for booking were anonymized. Prior to participation, all individuals received study information and were required to provide informed consent. Only participants who consented were included in the study, and data collection was limited to information necessary for the research.

5.6. Experimental Design

To ensure reliable outcomes, the experimental design integrated independent and dependent variables into a between-subject design [45], which exposes each participant to only version (Figure 6). This section details the independent and dependent variables examined, the recruitment, the tasks, and the data collection methods used.




Figure 6. Study design. The study employed one independent variable: Explanation type (Levels: Contrastive vs. Example-based), and four dependent variables: Comprehensibility, Simplicity, Cognitive Load, and User Experience. A total of 30 participants (15 per version) were recruited, with inclusion criteria requiring smartphone ownership and familiarity with Instagram.

5.7. Independent and Dependent Variables

The independent variable (IV) was the type of explanation, classified into two levels (L1: Contrastive explanation, L2: Example-based explanation).

The metrics used for the dependent variables (DVs) mentioned in Table 2 were collected in an Excel sheet and through online questionnaires created with Google Forms. The main outcomes of interest, and the methods for measuring each, are outlined below. All the metrics were recorded after participants interacted with the assigned prototype. The data collection units are as follows:

Table 2. Dependent variables and their measurement units.

Dependent Variables	Measurement (s)	Unit (s)
DV1: Comprehensibility	Mental model accuracy (1. objective and 2. subjective)	1. Rubric score (0–3); 2. Likert scale (1–7)
DV2: Simplicity	Irrelevant elements identified	Binary (yes/no)
DV3: Cognitive Load	NASA-TLX score	Likert scale (1–7)
DV4: User Experience	UEQ-S	7-Point (+3 to –3)

- DV1 (Comprehensibility): Objective comprehension was assessed through three open-ended questions targeting the 2W1H principle (what, how, and why) of data sharing privacy explanations. [16]. In addition, for subjective comprehension, a nine-item Likert scale questionnaire from the quality model (1–7 scale) was used.
- DV2 (Simplicity): Participants were asked to highlight any components or sentences in the explanation they considered irrelevant or unnecessary for their understanding.
- DV3 (Cognitive Load) was measured using a set of three Likert-scale items adapted from the NASA-TLX [40] questionnaire, assessing mental demand, effort, and frustration after the task.
- DV4 (User Experience; pragmatic, hedonic qualities) was assessed with the short version of the User Experience Questionnaire (UEQ-S) [41].

In addition to the DV, the participants also rated the explanations based on the adapted elements of aspects related to understandability, namely, effectiveness, trust, and transparency [8]. Participants were verbally asked to rate the explanations based on these aspects on a 7-level Likert scale. For completeness and transparency, the full list of user study questions has been included in Appendix B.

5.8. Study Procedure

As shown in Figure 7, the user study followed a sequential, three-phase procedure to ensure consistent data collection and minimize bias across participants. In the following, we will detail the procedure of the study.

5.8.1. Preparation

Step 1—Briefing and Demographics

Participants were briefed on the study's goal and provided with an overview of the tasks. Demographic questions were collected to describe the user sample.

Step 2—Baseline Alignment

Before showing the prototypes, all participants were shown a screenshot of how the existing Instagram location privacy notice appears on their device, to ensure that all participants had a similar shared baseline experience. This screenshot served only to give participants a common starting reference. It was not evaluated as a control or comparison condition, our analysis compares the two explanation formats with each other rather than against the existing notice.




Figure 7. Flow of the three-phase user study. Phase 1 included participant briefing and demographic questions. Phase 2 involved interacting with either the contrastive or example-based explanation prototype, followed by objective comprehension questions. Phase 3 consisted of four tasks: identifying unnecessary elements (simplicity), completing a subjective comprehension questionnaire, reporting cognitive load (NASA-TLX), and evaluating user experience (UEQ-S). Each step indicates approximate duration, whether responses were verbal or numeric, and the hypotheses addressed (H1–H4).

5.8.2. Prototype Interaction

Step 3—Explore Prototype

The prototype followed a single-user flow with the context of location data sharing explanations, embedded in the existing user interface of Instagram’s user feed, as detailed in Section 5.2. During this interaction, participants were explicitly instructed to focus solely on comprehending the provided explanation in terms of what is happening with their data (e.g., how their location data is being shared). They were not asked to perform any task during the interaction, such as allowing or denying location permissions based on the understanding, to isolate the measurement of understanding and avoid confounding variables like decision fatigue or external influences on behavior. While the overall prototype included interface elements such as toggles and checkboxes to support a realistic interaction flow, these elements were non-functional in the prototype.

Step 4—Evaluate Prototype

Immediately after interaction, we collected objective comprehension data by asking participants to answer open-ended questions about how the system works in terms of data sharing with advertisers. These questions were designed in accordance with the explanation provided to the participants in the interaction phase. It focuses on capturing the participant’s mental model accuracy and true understanding before any subjective impressions or biases from other questions could influence their responses.

Step 5—Additional Aspects

To align with the quality model’s recommendations and our research focus, we selected additional measurement items pertaining to effectiveness, trust, and transparency. Given that our aim was to assess how explanations impact understanding of data practices, these aspects were the most directly relevant. According to the model [8], effectiveness captures whether users feel empowered to make informed choices, trust reflects confidence in the privacy handling, and transparency ensures that the system provides enough information to the user to understand data practices, each of which is directly connected outcome of understandability in privacy interfaces. As these aspects were not the main focus of our study, we just took one item from each questionnaire to support our data. To evaluate these aspects, participants were required to rate these aspects of the explanations on a scale of 1–7 (1—Strongly disagree, and 7—Strongly agree).

5.8.3. Post-Interaction

This final phase measured all dependent variables as defined in Section 5.7, with each task designed to focus on a specific metric to measure the dependent variables.

Step 6—Tasks

Participants then completed four tasks in order. In Task 1, they identified any irrelevant or unnecessary elements in the explanation interface and provided open-ended feedback for each screen; although this feedback was not formally analyzed using the quality model, it was reviewed qualitatively to inform the design recommendations discussed in Section 6.4. In Task 2, participants filled out a nine-item Likert questionnaire (1–7), adapted from Deters et al. [8], to measure subjective comprehension part of the DV (mental model accuracy). Task 3 involved completing a three-item NASA-TLX [40]. Finally, in Task 4, participants completed the UEQ-S [41], an eight-item semantic differential questionnaire. The entire question list of the main study is provided in Appendix B.

5.9. Pilot Testing

A pilot test was conducted with four participants to validate the study design. The pilot confirmed an appropriate data distribution and produced no critical comments in the open feedback that would indicate problems with understandability. No issues were identified, so the procedure was retained for the main study.

5.10. Results—Exploratory Study

We recruited 30 participants (19 female, 11 male), aged 21–35 years ($M = 27.0$). Inclusion required smartphone ownership and regular Instagram use; individuals with professional experience in privacy research were excluded to avoid expert bias. All participants ($N = 30$) were currently university students. Participants reported high familiarity with the Instagram application ($M = 6.1$ on a 7-point scale). The sample consisted of 18 iOS (60%) and 12 Android (40%) users. We asked participants to mention the country they have resided for the most of their life, through which we found that the participant pool was culturally diverse. Participants were predominantly from India (43%, $n = 13$) and Germany (33%, $n = 10$). Other nationalities represented included Turkey (7%, $n = 2$), Bahrain (3%, $n = 1$), Pakistan (3%, $n = 1$), Bangladesh (3%, $n = 1$), Sri Lanka (3%, $n = 1$), and Indonesia (3%, $n = 1$). All of the participants were currently based in Germany, ensuring uniform exposure to local digital infrastructure and GDPR regulations [4].

5.10.1. Comprehensibility of Presentation Form

We tested the difference between the two versions using objective metrics and subjective scores. The objective metric involved rubric-scored responses to open-ended questions, while the subjective metric drew from Likert-scale items assessing perceived understanding. Followed by a Spearman's rho correlation analysis to explore any relationship between them. This dual approach allowed for a comprehensive assessment, capturing both actual knowledge demonstrated through responses and self-reported perceptions of comprehension.

Objective Comprehension

We analyzed the responses for each participant using a scoring rubric that we created to structure participant responses, enabling rubric-based scoring (0–3 scale) that quantifies objective accuracy without relying on self-reported perceptions, which could introduce subjectivity. The scoring Rubric was designed for each question response. We assigned 1 point for accurate, complete responses aligned with the system's actual mechanics (e.g., mentioning city-level or anonymized data without exact GPS); 0.5 points for partially correct but

ambiguous answers (e.g., vague references to “areas” or “aggregation” without clarity); and 0 points for incorrect or absent understanding (e.g., assuming real-time tracking or individual identification). This granularity ensured the rubric reflected varying levels of insight, with a maximum total of 3 points per participant. While the rubric enabled structured and consistent evaluation, it was developed specifically for this study and may reflect some level of subjective interpretation by the author. Descriptive statistics (Table 3) indicated slightly better performance (in terms of objective comprehension) with example-based explanations (Interface B: $M = 2.4$, $SD = 0.60$, $Md = 2.5$) than with contrastive explanations (Interface A: $M = 2.03$, $SD = 0.74$, $Md = 2.0$). Box plot visualizations (Figure 8) showed that Interface B not only yielded higher mean scores but also less variability, suggesting more consistent mental models. Given the ordinal nature of the rubric scores, we applied a Mann–Whitney U test to compare groups, as non-parametric methods are recommended for ordinal outcomes without assuming normality or equal intervals [46]. A Mann–Whitney U test found no statistically significant difference between the two versions ($U = 79.5$, $z = -1.35$, $p = 0.177$).

Table 3. Descriptive statistics for mental model accuracy (Objective scores).

Metric	Interface A	Interface B
Mean	2.033	2.4
Median	2.0	2.5
Std. Dev.	0.743	0.604




Figure 8. Box plot visualizing objective comprehension scores (rubric scale: 0–3) for Interface A and Interface B. Interface B displays a slightly higher mean ($M = 2.4$) and median ($Md = 2.5$) compared to Interface A ($M = 2.033$, median = 2.0), indicating somewhat better performance. The shorter interquartile range (IQR) and whiskers in Interface B suggest reduced variability and more consistent comprehension across participants.

Subjective Comprehension

We measured subjective comprehension with a 9-item Likert questionnaire (7-point scale) adapted from the explainability quality model [8]. Originally, the scale included 10 items (Table 4). We adapted the items to fit our between-subjects design: one comparative item (1.5) was removed, and two items (1.4 confusing terms, 1.6 mental demand) were reverse-coded to align scoring so that higher values consistently indicated more positive perceptions. Composite comprehension scores were calculated as the mean of all items per participant.

Both interfaces achieved high perceived comprehension scores, well above the scale midpoint of 4 (Table 5). Interface A (contrastive) scored $M = 5.7$ ($Md = 6.1$, $SD = 0.7$),

while Interface B (example-based) scored $M = 5.8$ ($Md = 6.0$, $SD = 0.7$). Box plot inspection (Figure 9) showed nearly identical distributions, with medians close together and similar variability. These patterns suggest that participants generally viewed both interfaces as easy to understand.

Table 4. Questionnaire items measuring subjective perception for understandability aspect [8].

ID	Question
Q1.1	I found it easy to understand why/how the system did [event].
Q1.2	The explanations provided made sense to me.
Q1.3	Some explanations did not make sense to me, because ...
Q1.4	The explanation contains terms that are confusing to me.
Q1.5	I think the explanations from system A can better help me understand [event, result] compared to system B. (excluded from the actual experiment)
Q1.6	I found the task very mentally demanding.
Q1.7	I find the explanation easy to read.
Q1.8	I find the length of the explanation appropriate.
Q1.9	I find that the explanation is written in correct, appropriate English.
Q1.10	The arrangement/organization of information appears very logical to me.

Table 5. Descriptive statistics for mental model accuracy (Subjective scores).

Metric	Interface A	Interface B
Mean	5.7	5.8
Median	6.1	6.0
Standard Deviation	0.7	0.7




Figure 9. Box plot comparing subjective comprehension scores (7-point Likert scale) between Interface A and Interface B. Both interfaces show similarly high scores above the scale's midpoint of 4, with closely aligned medians (6.1 for A, 6.0 for B) and identical standard deviations ($SD = 0.7$). The similar spread of the boxes and whiskers indicates comparable variability, supporting the interpretation that both explanation types were perceived favorably and yielded similar subjective understanding across participants.

Given the evidence of non-normality as indicated by the Shapiro-Wilk test in Interface A ($W = 0.8757$, $p = 0.041$), we applied a Mann–Whitney U test. The test showed no significant difference between groups ($U = 107$, $z = -0.23$, $p = 0.833$). Together, the descriptive and inferential results indicate that both types of explanation were perceived as highly comprehensible. Because subjective scores were non-normal and compared with a Mann–Whitney U test, we describe central tendency using the median rather than the mean. By median, Interface A ($Md = 6.1$) and Interface B ($Md = 6.0$) were nearly identical,

with Interface A marginally higher. As the difference was not statistically significant, we do not interpret a directional trend in subjective comprehension.

5.10.2. Correlation

In addition, to examine the relationship between objective performance and subjective perceptions of understanding, we conducted Spearman's rank correlation analysis [46]. We compared participants' actual comprehension scores with their self-reported understanding to examine whether people can accurately judge their own comprehension. Prior research shows that people often overestimate or misjudge their understanding, creating a gap between perceived understanding and actual performance [47].

For Interface A, the correlation between participants' actual comprehension scores and their self-reported understanding was weakly positive ($\rho = 0.077$), but this association did not reach statistical significance ($p = 0.784$). Similarly, Interface B exhibited a weak negative correlation ($\rho = -0.102$) that was also non-significant ($p = 0.719$). These findings suggest that, regardless of interface design, participants' confidence in their understanding did not reliably correspond to their demonstrated task performance. In other words, participants were generally unable to accurately assess their own comprehension, indicating a disconnect between perceived and actual understanding across both experimental conditions. This finding reflects the distinction between perceived and objective understanding within our overall study design. In the pre-study (c.f., Section 4), participants rated only their perceived understanding of Smartphone data practices, which captures familiarity and confidence rather than accuracy. In the main study, we measured both subjective comprehension (participants' perceptions of how understandable the explanations were) and objective comprehension (their demonstrated mental-model accuracy). The weak and non-significant correlations show that these two forms of understanding do not align: participants often felt that they understood the explanations even when their mental-model representations were incomplete or incorrect.

5.10.3. Simplicity

Simplicity was measured by asking participants to identify whether any elements of the explanation were unnecessary for helping them understand the content. Participants provided a binary response ("Yes" or "No"), allowing us to calculate the proportion of users in each interface who perceived irrelevant elements in the explanation.

Most participants (80.0%, $N = 24$) found no irrelevant elements, indicating that both interfaces were generally perceived as simple to understand. As shown in Table 6, in Interface A (contrastive), 13.3% flagged unnecessary content, compared to 26.7% in Interface B (example-based). While more participants in Interface B reported the occurrence of irrelevant content for their understanding, the majority of participants in both groups did not.

Table 6. Distribution of irrelevant elements identified ("Yes" or "No").

Group	Y (Count)	Y (%)	N (Count)	N (%)	Total
Interface A	2	13.3%	13	86.7%	15
Interface B	4	26.7%	11	73.3%	15
Total	6	20.0%	24	80.0%	30

Given the binary categorical nature of the dependent variable and the independent groups design, a Chi-square test of independence was selected [46]. A chi-square test of independence confirmed no statistically significant difference ($\chi^2 = 0.83$, $p = 0.361$).

While our primary measure of simplicity was the binary response indicating whether participants perceived any element as unnecessary, we also collected qualitative insights by asking participants to briefly explain why they considered a specific part of the explanation unnecessary. This allowed us to capture user feedback and identify areas for improving the explanation designs, in particular in terms of simplicity. To interpret the qualitative feedback on irrelevant or unnecessary elements that users pointed out on the screens, we conducted a close reading of the comments. A researcher reviewed all responses and noted recurring observations related to redundancy, clarity of wording, or information density. These observations were then summarized into interface-level patterns without applying a formal coding scheme, reflecting the small scale and targeted nature of the qualitative data. Among the minority who identified irrelevant elements, distinct patterns emerged. In Interface A, two participants (P17 and P25) felt the section describing “what is not shared” was redundant, as they already understood the explanation from the “what is shared” content alone. For Interface B, four participants (P04, P08, P14, P24) highlighted the “if denied” scenario as unnecessary, noting that the “allowed” example was sufficient for comprehension. One participant also described the example text as “heavy and wordy,” suggesting that two-scenario examples made the explanation feel denser.

Although no significant differences were found, qualitative feedback points to subtle design tensions: contrastive explanations may risk redundancy by emphasizing “what is not shared,” while example-based explanations may feel wordy when including both allowed and denied scenarios. Together, these findings suggest that tailoring explanations to highlight the most relevant information (what is shared or allowed) may improve simplicity without sacrificing informativeness.

In addition to feedback on the explanation screens, several participants expressed confusion about the lock icon used on Screen 1 to initiate the flow. Although this element was not evaluated as a dependent variable, it formed part of the interaction sequence, which resulted in spontaneous comments from participants. Several participants misunderstood icons such as the lock icon in screen 1 (cf. Appendix C), which was used as a trigger in our case. In total, 60% of participants (18 out of 30 participants) raised concerns or confusion regarding the lock icon’s relevance. P05 stated “The lock icon looked like an shopping bag to me, and first I thought it was relevant for the advertisement and not privacy”. While a few participants noted its visibility or familiarity with privacy settings from other apps, many criticized its lack of clarity, contextual fit, and clickability. These findings reinforce the need for pairing icons with descriptive text to reduce ambiguity in privacy-related UI elements.

5.10.4. Cognitive Load

Cognitive load was measured using a subset of the NASA-TLX [40]. In the analysis, only three sub-scales were selected: (1). Mental Demand (“How mentally demanding was it to understand the explanations?”), (2). Effort (“How hard did you have to work to fully understand the explanations?”), and (3). Frustration (“How frustrated did you feel while reading or interacting with the explanations?”), as they directly captured the cognitive and emotional strain relevant to processing privacy explanations in a non-physical task. The other sub-scales were excluded because Physical demand does not apply to a primarily screen-based interaction without motor/manual requirements; Temporal demand is irrelevant as there were no time pressures in the study protocol; and Performance focuses on self-evaluated success rather than comprehension. This targeted adaptation maintained the tool’s integrity while aligning with the research focus on mental load in explainable privacy interfaces. Participants rated each item on a 7-point Likert scale, and composite scores were calculated by averaging the three items.

Both interfaces were associated with low cognitive load, with means below the midpoint of 4. As shown in Table 7 Interface A (contrastive) scored slightly higher ($M = 2.80$, $SD = 1.52$) than Interface B (example-based; $M = 2.38$, $SD = 1.15$). These results suggest that participants found both types of explanation manageable in terms of cognitive load, with a minor trend toward higher cognitive strain in the contrastive condition.

Table 7. Descriptive statistics of perceived cognitive load scores across interfaces.

Interface Version	Mean (M)	Standard Deviation (SD)
A (Contrastive)	2.80	1.52
B (Example-based)	2.38	1.15

Assumptions of normality were met, allowing for parametric comparison. An independent-samples t-test revealed no significant difference between the two interfaces, $t(28) = 0.86$, $p = 0.399$.

While descriptive trends indicated that contrastive explanations may impose slightly higher mental effort, the difference was not statistically significant.

5.10.5. User Experience

To complement the core measures of comprehensibility, simplicity, and cognitive load, we assessed overall user experience of both interface versions using the UEQ-S [41].

Both interfaces performed strongly in terms of pragmatic quality but weaker in terms of hedonic quality (Table 8). For Interface A, pragmatic quality averaged 1.85 ($SD = 0.71$, “excellent”), hedonic quality 0.83 ($SD = 0.91$, “below average”), and overall 1.34 ($SD = 0.65$, “good”). Interface B produced nearly identical results: pragmatic $M = 1.87$ ($SD = 0.68$), hedonic $M = 0.83$ ($SD = 0.95$), overall $M = 1.35$ ($SD = 0.63$). Benchmark comparisons classified pragmatic scores as excellent, hedonic as below average, and overall experience as good.

Table 8. User experience mean and standard deviation (UEQ-S).

Quality Type	Mean	SD
Interface A (Contrastive)		
Pragmatic Quality	1.85	0.712
Hedonic Quality	0.83	0.910
Overall	1.34	0.649
Interface B (Example-based)		
Pragmatic Quality	1.87	0.681
Hedonic Quality	0.83	0.948
Overall	1.35	0.629

An independent samples t-test was conducted to compare the groups, as the normality assumptions and equal variances were met. Independent-samples t-tests confirmed no significant differences between interfaces for pragmatic ($t(28) = -0.07$, $p = 0.948$) or hedonic quality ($t(28) = 0.00$, $p = 1.000$). These results indicate no significant difference. These findings suggest that both explanation types were perceived as highly usable and supportive but lacking in excitement.

5.10.6. Additional Aspects Related to Understandability

As mentioned in Section 5.6, our analysis targeted three additional key aspects from the quality model: Effectiveness, Trustability, and Transparency [8]. Quantitative analysis revealed that Interface A (contrastive) consistently received higher average ratings than

Interface B (example-based): Effectiveness ($M_A = 6.0$, $M_B = 4.5$); Trustability ($M_A = 5.5$, $M_B = 4.5$); Transparency ($M_A = 6.5$, $M_B = 6.0$).

These descriptive results suggest that the contrastive explanation interface (Interface A) more effectively communicated data-sharing practices to participants across all three items. However, these differences were not statistically tested and should therefore not be interpreted as evidence of superiority. Instead, they provide preliminary indications of how participants perceived the two explanation types, which we discuss as exploratory patterns.

6. Discussion

This chapter brings together the results to address the research questions introduced at the start of the paper. Its purpose is to examine whether the hypotheses can be supported or rejected based on the collected evidence. For this, insights from the literature review, the pre-study survey with 100 participants, and the user study ($N = 30$) are considered jointly. The discussion highlights the main conclusions drawn from these findings.

The central research question asked: How can explainable privacy interfaces be designed to enhance user understanding of app-specific data practices?

We address this through insights from the three sub-research questions, which together inform our final design recommendations.

6.1. Sub-RQ1: Users Struggle Most with Third-Party Sharing but Show High Interest

Our first research question asked: What is the current level of user understanding and interest regarding app-specific data practices (collection, usage, and third-party sharing)?

Consistent with prior work on privacy policies and runtime notices [6,7], our pre-study survey ($N = 100$) revealed uneven comprehension. More than half of participants (51%) reported that they mostly or fully understood data collection, while fewer did so for data usage (40%) and especially for third-party sharing (30%). Conversely, 44% indicated that they did not understand or only slightly understood data sharing, compared to 30% for data usage and 20% for data collection. This highlights that sharing practices remain especially opaque, aligning with prior literature that points to hidden, legalistic, or fragmented disclosures [3,22,23].

Despite these comprehension gaps, participants expressed a strong interest in learning more. Three-quarters (76%) reported being mostly or fully interested in further information about data sharing, with slightly lower but still high levels of interest in data usage (69%) and collection (67%). This pattern not only mirrors prior findings that users are concerned about hidden third-party flows [9,24], but also suggests a gap between current understanding and user desire for transparency. Our findings therefore support the view that future privacy interfaces should particularly target explanations of third-party sharing, as comprehension is lowest but interest is highest for third-party data sharing practices.

6.2. Sub-RQ2: Privacy Explanations Help, but Fall Short on Sharing

Our second research question asked: How do users perceive current privacy explanations, and what improvements do they suggest?

Consistent with prior work on privacy explanations [16], our survey showed strong demand for explanations in app interactions. When asked how important it is that an app explains why it requests location access, 63% of participants rated this as “Very important” or “Fully important,” with only 3% indicating “Not at all important.” This confirms that users see explanations as a critical component of transparency.

When provided with an adapted privacy explanation, participants found it helpful for clarifying access permissions. For allowing access, 64% rated the explanation “Very” or “Fully helpful,” and for denying access, 58% did so. This aligns with findings of

Brunotte et al. [16], who emphasize that explanations enhance awareness and trust. However, third-party sharing remained problematic: only 44% rated the explanation as “Very” or “Fully helpful,” while 37% found it “Not helpful” or only “Slightly helpful.” This confirms that current formats still struggle to make third-party practices comprehensible.

Open-text responses reinforce these findings. Participants called for greater transparency about what data is shared, why, and with whom ($N = 43$). Others emphasized the need for visual aids ($N = 40$). Additional suggestions included shorter contextual copy with progressive disclosure ($N = 16$), assurances of legal compliance ($N = 6$), and limited personalization options ($N = 2$). These themes resonate with prior work on contextual privacy policies [9], but highlight that even when explanations are present, users still struggle most with third-party sharing.

Together, these results suggest that while explanations support comprehension of allow/deny decisions, they fall short for third-party practices. Addressing this gap will require designs that combine transparency, visual support, layers, and contextual delivery, as reflected in our derived design guidelines (Table 1).

6.3. Sub-RQ3: Example-Based and Contrastive Explanations Show Complementary Strengths

Our third research question asked: How do different explainable privacy interfaces (example-based vs. contrastive) influence user understanding of app-specific data practices?

To address the gap identified by our RQ2, we explored whether more advanced explanation strategies could be applied. In the explainable AI literature, contrastive and example-based explanations are widely studied as two influential approaches for improving user understanding [3,30,31]. Unlike regular textual explanations, these techniques emphasize either clarifying why not alternatives (contrastive) or illustrating outcomes through relatable cases (example-based) [3]. Building on user feedback from RQ2, we designed and tested both types of explanations in our user study with 30 participants to investigate whether they can better support comprehension of app-specific data practices (RQ3).

Before interpreting these results, we reiterate that this is an exploratory study. With 30 participants (15 per condition), so the comparisons below should be read as descriptive trends that motivate future, larger-scale work rather than as confirmed differences between the two explanation types.

Across comprehension, simplicity, cognitive load, and user experience, inferential tests showed no statistically significant differences between explanation types. This suggests that both versions were comparably effective. However, descriptive trends of our study revealed complementary strengths.

Example-based explanations showed slightly higher objective comprehension scores, which aligns with Adadi and Berrada’s view that example-driven approaches support user understanding by grounding system behavior in concrete instances [33]. Contrastive explanations received marginally higher ratings on the effectiveness, trustability, and transparency, consistent with Miller’s observation that human reasoning is often contrastive and oriented toward distinguishing why one outcome occurs rather than another [32]. Such contrastive framing may help users form clearer boundaries about what data is and is not shared and may therefore feel more supportive in decision-oriented contexts, as reflected in the additional aspect’s effectiveness rating. These interpretations remain exploratory, since the study was not powered to detect small differences, but the findings suggest that example-based explanations may better support comprehension, while contrastive explanations may play a complementary role in decisions involving privacy-related choices.

These insights suggest explanation design should be purpose-driven: example-based explanations should be prioritized when the goal is to provide information and comprehen-

sion, and contrastive explanations should be used when the goal is to aid user decisions (Figure 10). In contexts where both are needed, progressive disclosure strategies may balance clarity and cognitive load.

In addition, we formulated four hypotheses (H1–H4) in Sub-RQ3 to examine differences between the two interface variants (Interface A: contrastive and Interface B: example-based) with respect to specific DVs (cf. Table 2). Below, we summarize and discuss the results of this evaluation.




Figure 10. Aligning explanation type with user-intended purpose.

6.3.1. “There Is a Significant Difference in User Comprehension Between Interface A and Interface B, as Measured by Mental Model Accuracy” (H1)

Descriptive statistics indicated slightly higher mental model accuracy for example-based explanations ($M_B = 2.4$) compared to contrastive ($M_A = 2.03$), with less variability suggesting more consistent comprehension. However, both objective comprehension scores (Mann–Whitney $U = 79.5$, $z = -1.35$, $p = 0.177$) and subjective comprehension ratings (Mann–Whitney $U = 107$, $z = -0.23$, $p = 0.833$) revealed no significant differences between the two versions. Therefore, H1 was not supported. This suggests that both explanation types facilitated a high level of understandability corroborating prior findings that layered, contextual privacy explanations enhance clarity [9,24]. Despite non-significant differences, example-based explanations may modestly improve mental model consistency by providing concrete counterfactuals [33].

6.3.2. “There Is a Significant Difference in Perceived Simplicity Between Interface A and Interface B, as Measured by the Identification of Irrelevant Elements in the Explanation” (H2)

Descriptive trends revealed that the majority (80%) of participants found no irrelevant elements in either explanation. A chi-square test on binary relevance responses confirmed no significant difference between conditions ($\chi^2 = 0.83$, $p = 0.361$). Therefore, H2 was not supported.

Consequently, both explanation styles demonstrate the capability to distill complex privacy concepts into accessible forms, making them suitable candidates for enhancing user understanding in app-specific data-sharing contexts. While qualitative feedback highlighted minor stylistic tensions: contrastive explanations risk redundancy (e.g., “what is not shared”), and while example-based explanations may feel wordy due to multiple

scenarios, these did not detract from the overall clarity or simplicity perceived by most users. Therefore, such explanation approaches can be trusted to facilitate comprehension and engagement, provided designers continue to focus on maintaining simplicity, which is a crucial criterion in the understandability aspect of explanations [8].

6.3.3. “There Is a Significant Difference in Perceived Cognitive Load Between Interface A and Interface B, as Measured by the NASA-TLX Scale” (H3)

Both interfaces imposed a low cognitive load (mean scores below the midpoint of 4), with contrastive explanations scoring slightly higher ($M = 2.80$) than example-based explanations ($M = 2.38$). However, results from an independent-samples t-test showed no significant difference ($t(28) = 0.86, p = 0.399$). Therefore, H3 was not supported. These findings indicate that cognitive demands are manageable for both explanation types, with minor trends suggesting contrastive explanations incur slightly higher mental effort, possibly due to processing “not shared” information as outlined above in the interpretation of H2 results.

6.3.4. “There Is a Significant Difference in User Experience Between Interface A and Interface B, as Measured by the User Experience Questionnaire (UEQ-S)” (H4)

User experience scores indicated excellent pragmatic quality but below-average hedonic quality for both interfaces. An independent-samples t-test showed no significant differences for pragmatic quality ($t(28) = -0.07, p = 0.948$) or hedonic quality ($t(28) = 0.00, p = 1.000$). Therefore, H4 was not supported. This suggests both explanation styles are usable and supportive, but lack engaging or enjoyable elements. This pattern aligns with prior research on engagement in privacy notices, which highlights that privacy notices often fail to capture sustained user interest or emotional engagement [7]. In light of these findings, our results reinforce the need for privacy explanation interface designers to go beyond basic interaction methods: designers should consider incorporating features that enhance engagement, such as interactive elements like Swipe and Drag-and-Drop (DAD) actions [7], to overcome the limitations of low hedonic quality and prevent user disengagement in privacy-relevant contexts.

6.4. Recommendations to App Designers

Drawing on the findings from our study and the broader privacy and explainability literature, we derive the following recommendations for the design of explainable privacy interfaces in smartphone applications to answer our central research question: How can explainable privacy interfaces be designed to enhance user understanding of app-specific data practices?

- Recommendation 1: Use clear, labeled triggers. As discussed in Section 5.10.3, qualitative feedback showed that icons alone were often insufficient for conveying privacy-related meaning, leading to misunderstandings of their function. Pairing icons with short, descriptive text reduces ambiguity and supports user recognition of privacy explanations. In practice, this confusion could be reduced by pairing the trigger icon with a short micro-text label (for example, ‘Privacy details’), so its function is not left to interpretation. Animated affordances or a brief first-use tooltip could further signal that the element is interactive and privacy-related rather than decorative. Such lightweight cues address the misrecognition we observed. This is consistent with prior evidence that unclear or overly brief cues hinder user comprehension of data practices [6,7]. Designers should therefore rely on explicit labeling alongside visual cues to increase clarity.
- Recommendation 2: Condense data granularity through progressive disclosure. As detailed in the qualitative findings in Section 5.10.3 about simplicity of explanations,

revealed that listing fine-grained details of data can impose a cognitive burden on participants, which can dilute the central message about data practices. Grouping such details under meaningful higher-level categories (e.g., presenting “location data” as a single conceptual unit rather than separating allowed/denied sub-components) can reduce perceived complexity. Progressive disclosure may further support this by allowing users to access additional detail only when needed. This aligns with calls in the literature to reduce information overload and ensure contextual relevance of notices [9,24]. Designers should therefore present privacy explanations in a contextual, step-by-step layered format (e.g., what the entities are, why sharing occurs, and how); these findings are in line with Brunotte et al. [16]’s 2W1H principle for privacy explanations.

- Recommendation 3: Do not rely solely on self-assessment for understanding. We found that participant confidence in understanding explanations did not always correlate with actual comprehension, echoing findings that users often misinterpret runtime privacy notices [6]. Instead of depending on self-reports, privacy interfaces should embed lightweight feedback mechanisms, such as interactive comprehension checks, to ensure that information is not only accessed but also meaningfully understood.
- Recommendation 4: Provide granular third-party sharing controls with tailored explanations. The most persistent comprehension gap concerned third-party data sharing, which participants identified as both confusing and critically important. This is in line with previous findings that third-party sharing represents a major knowledge gap where vague or absent explanations fail to inform users [21,23,48]. In our study, participants appreciated the fine-grained, revocable controls that differentiated between categories of third parties to help them. This observation highlights the work by Schaub et al. [23], who demonstrated that users frequently struggle to comprehend third-party descriptions in privacy notices when these are limited to company names or vague categories, such as “data reseller” or “government agency.” Hence, it is necessary to provide users with granular, explanatory controls that clarify not only who third parties are but also how and why their data is shared.

6.5. Limitations and Threats to Validity

This study has several limitations. To assess the sensitivity of the study, we conducted a post-hoc sensitivity analysis using G*Power (version 3.1). For $N = 30$, $\alpha = 0.05$, and a desired power of 0.80, the study was powered to detect only large effects (Cohen’s $d \approx 1.06$). Smaller effects may therefore have remained undetected, and statistical power should be considered a limitation of the study. Moreover, the participant pool lacked demographic diversity, being largely young and digitally literate, which may obscure the privacy needs of older adults or less tech-savvy users. The study further considered only end-user perspectives, excluding developers whose technical and organizational constraints critically shape the feasibility of explainable privacy features; future work should integrate their insights through co-design or participatory approaches. Our evaluation focused on a single location-sharing scenario. While this reflects a common and relatable permission request, it limits the generalizability of the findings to other app contexts and data types. Future work should test multiple scenarios to determine whether the observed patterns hold across different privacy situations. Finally, our evaluation relied on the explainability quality model by Deters et al. [8], which was developed primarily in other domains such as recommender systems, and its applicability to privacy contexts remains preliminary.

7. Conclusions

In this paper, we investigated how explainable privacy interfaces can enhance user understanding of app-specific data practices. A pre-study survey with 100 participants, followed by a comparative user study ($N = 30$), revealed that users experience the greatest difficulty in understanding third-party data sharing, despite showing strong interest in this area. Building on this, we extended prior knowledge by testing two explainability strategies, contrastive and example-based explanations, in the context of the understandability of privacy interface explanations. While we found no statistically significant differences across comprehension, simplicity, cognitive load, or user experience, descriptive results revealed complementary strengths. We position this work as exploratory research. Its findings derive from a small sample ($N = 30$, 15 per condition) without a conventional-notice baseline and did not reach statistical significance, so they should be interpreted as promising descriptive trends rather than confirmatory effects. Their value lies in identifying directions, such as matching explanation type to user goal, that motivate future work with larger, adequately powered, and baseline-controlled studies.

7.1. Contribution

Our findings contribute actionable implications for the design of explainable privacy interfaces: explanations should be layered, concise, contextual, and visual, with particular emphasis on transparency in third-party sharing. Designers should further align explanation styles with user goals and intended purpose, using example-based approaches for comprehension and contrastive ones for decision support. These insights can guide app developers in creating privacy explanations that move beyond regulatory compliance towards genuine user understanding of data practices.

7.2. Future Work

This work focused primarily on understandability as a quality aspect of explainable privacy interfaces. Future research should extend evaluation to the other aspects defined in the quality model by Deters et al. [8]. Considering the dimensions in combination could reveal trade-offs or synergies, leading to a more holistic understanding of how explanations can affect user understanding of privacy. Another promising direction is the integration of explainable AI (XAI) features directly into privacy interfaces. While our prototypes employed static designs, adaptive explanation generation, user-tailored feedback, or intelligent privacy assistants could provide more dynamic support. Evaluating such AI-driven features in real application contexts would help assess their potential for delivering real-time, personalized, and trustworthy privacy explanations.

Author Contributions: Conceptualization, A.B. and A.R.; methodology, A.B. and A.R.; software, A.B.; validation, A.B. and A.R.; resources, A.R.; data curation, A.B.; writing—original draft preparation, A.B. and A.R.; writing—review and editing, A.B. and A.R.; visualization, A.B.; supervision, A.R.; project administration, A.B. and A.R.; funding acquisition, A.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Participation was in accordance with the ethical guidelines stated in the Declaration of Helsinki: participation was voluntary, participants were obliged to provide their written informed consent, and had the opportunity to abort the study at any time and leave comments or feedback. Study participants were not exposed to any kind of physical or psychological harm during the study; therefore, based on our university's regulations, ethical approval was not required.

Informed Consent Statement: Informed consent was obtained from all subjects prior to the beginning of the study.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: In this work, AI tools were used to correct grammatical errors, paraphrase sentences, search for related works, and make stylistic improvements. We carefully made these adjustments, considering the original content, with the aim of enhancing the quality and clarity of this work. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Pre-Study Survey Questionnaire

Section 1: Demographics

1. What is your age?
 - Under 18
 - 18–24
 - 25–34
 - 35–44
 - 45–54
 - 55–64
 - 65 or older
2. What is the gender you most identify with?
 - Male
 - Female
 - No answer
 - Other
3. What is your current profession?
 - Student
 - Academic or Researcher
 - Professional (e.g., Engineer, IT Specialist, Business Analyst)
 - Educator or Teacher
 - Healthcare Professional
 - Artist or Designer
 - Legal Professional
 - Environmental Specialist
 - No answer
 - Other
4. What is your field of studies?
 - Computer Science/Information Technology
 - Engineering
 - Business/Economics/Management
 - Social Sciences
 - Natural Sciences
 - Health Sciences
 - Humanities
 - Arts and Design
 - Law
 - Education

- Environmental Sciences
- No answer
- Other

Section 2: General Data about Smartphone Usage

1. Do you own a personal Smartphone?
 - Yes
 - No
2. What is the language setting of your Smartphone?
 - English
 - German
 - No answer
 - Other
3. What is your Smartphone's primary operating system?
 - iOS
 - Android
4. What type of Smartphone applications do you use on a daily basis? (Select all that apply)
 - Social Media
 - Streaming and Entertainment
 - Productivity and Work
 - Fitness and Health
 - Shopping and E-commerce
 - Banking and Finance
 - Travel and Navigation
 - News and Information
 - Gaming Apps
 - Education and Learning
 - No answer
 - Other

Section 3: Understanding of App-Specific Privacy Settings

1. Are you aware that Smartphone applications collect data about you (e.g., location data)?
 - Not at all aware
 - Slightly aware
 - Moderately aware
 - Mostly aware
 - Fully aware
2. How well do you understand the purpose of the following settings? (Likert scale: 1 = Do not understand at all, 5 = Fully understand)
 - Device Access
 - Security
 - Internet Connectivity
 - Location
 - Notifications
3. How well do you understand the collect, use, and share of your personal data through these settings? (Likert scale: 1 = Do not understand at all, 5 = Fully understand) (Please choose the response for each item)

- How apps collect your data
- How apps use your data
- How apps share your data

Note—Data Collection: The process by which an application acquires user information.

Note—Data Usage: The ways in which collected information is internally processed.

Data Sharing: The disclosure or transfer of user information to external entities, which may include third-party services.

- Are you interested in learning about how your personal data is COLLECTED through Smartphone applications?

- Not at all interested
- Slightly interested
- Moderately interested
- Mostly interested
- Fully interested

Note—Data Collection: The process by which an application acquires user information.

- Which settings would you like to learn more about regarding data COLLECTION? (Select all that apply)

- Device Access
- Security
- Internet Connectivity
- Location
- Notifications
- All of the above
- None
- Other

- Are you interested in learning about how your personal data is USED by Smartphone applications?

- Not at all interested
- Slightly interested
- Moderately interested
- Mostly interested
- Fully interested

Note—Data Usage: The ways in which collected information is internally processed.

- Which settings would you like to learn more about regarding data USAGE?

- Device Access
- Security
- Internet Connectivity
- Location
- Notifications
- All of the above
- None
- Other

- Are you interested in learning about how your personal data is SHARED by Smartphone applications?

- Not at all interested
- Slightly interested

- Moderately interested
- Mostly interested
- Fully interested

Note—Data Sharing: The disclosure or transfer of user information to external entities, which may include third-party services.

9. Which settings would you like to learn more about regarding data SHARING?

- Device Access
- Security
- Internet Connectivity
- Location
- Notifications
- All of the above
- None
- Other

Note—Data Sharing: The disclosure or transfer of user information to external entities, which may include third-party services.

Section 4: Effectiveness of Privacy Explanations

1. Please read the following scenario and then answer the question:

Scenario: You have downloaded a new smartphone app called “CityExplorer” that helps users discover hidden attractions, local restaurants, and personalized sightseeing routes. When you open “CityExplorer” for the first time, it requests access to your Location without providing any explanation.

How important is it that an app tells you why it needs to use your location?

- Not at all important
- Slightly important
- Moderately important
- Very important
- Fully important

2. “Now, imagine the app provides the following explanations”:

- Implications of allowing access: Enhances user experience by providing you personalized recommendations.

Note: Location data may be shared with advertisers.

- Implications of denying access: You get generic recommendations.

Note: Location data will not be shared with advertisers.

How helpful are the explanations in helping you understand the following? (Likert scale: 1 = Not helpful at all–5 = Fully helpful) (Please choose the response for each item)

- Implications of allowing access
- Implications of denying access
- How data is shared with third parties

3. What do you think about the level of detail in the explanation?

- Not detailed at all
- Slightly detailed
- Moderately detailed
- Mostly detailed
- Fully detailed

4. During what time would you find these explanations helpful?

- Before app installation or during setup
 - At the moment of permission request
 - Upon first use of a feature requiring permissions
 - Periodically, as reminders
 - On-demand
 - I don't find these explanations helpful
 - Other
5. How did the explanation affect your understanding or perception of the app's use of your data? (Open text)
 6. What do you think should be improved about the explanation? (Open text)

Appendix B. Main Study Questionnaire

Section 1: Demographics (Collected Verbally)

- Participant ID
- Version (A/B)
- Age
- Gender
- Country where the participant has resided for the majority of their life
- Profession
- Device OS (iOS/Android)
- Comfort level in using Instagram (1–7)

Section 2: Post-Prototype Interaction (Collected Verbally)

1. Location data mental model (open-ended):
 - (a) What type of location data do you think is shared with advertisers? (e.g., exact GPS, city-level location, tagged places)
 - (b) How do you think advertisers see your location data? (e.g., grouped into categories)
 - (c) Why do you think this location data is shared with advertisers? (e.g., to show nearby ads, improve app performance, track movements)
2. Effectiveness:
"I find the explanation provided has sufficient information to make an informed decision to allow/deny the location sharing." (1–7 Likert)
3. Trustability:
"The explanations increased my trust in the application with respect to my location data privacy." (1–7 Likert)
4. Transparency:
"Being able to see and control exactly what types of data are shared about me (e.g., age range, interest categories) helps me better understand the app's data practices." (1–7 Likert)

Section 3: Screen 4—Explanation Screens

Participants verbally completed the following questions for Screen 4:

1. Screen 4: Anything irrelevant? (Y/N)
Remove elements (e.g., sentences) from the explanation that were not necessary for your understanding.
(Focus on the text explanation)
2. Screen 4: Which part?
(Open-ended response)

3. Screen 4: Can you specify the reason?
(Open-ended response)
4. Screen 4: Suggest any change in the specific part which will help your understanding
(Open-ended response)
5. Screen 4: Simplicity (1–7). Why?
(Open-ended explanation for the rating)

Section 4: Subjective Comprehension

Participants rated the explanations on 7-point Likert scales (1 = Strongly Disagree, 7 = Strongly Agree):

1. I found it easy to understand why/how the system showed the ad based on my location.
2. The explanations provided made sense to me.
3. The explanations contain terms that are confusing to me.
4. The explanations felt logically structured and easy to follow.
5. I found the overall task very mentally demanding.
6. I find the explanations easy to read.
7. I find the length of the explanations appropriate.
8. I find that the explanations are written in correct, appropriate English.
9. The arrangement/organization of information in the explanations appears very logical to me.

Section 5: NASA-TLX

Participants assessed the cognitive load of the explanations using three NASA-TLX subscales (1 = Very Low, 7 = Very High). The items were presented exactly as follows:

1. Mental Demand:
How mentally demanding was it to understand the explanations?
2. Effort:
How hard did you have to work to understand the explanations?
3. Frustration:
How frustrated did you feel while reading or interacting with the explanations?

Section 6: User Experience (UEQ-S)

Participants rated each of the following adjective pairs on a 7-point likert scale, using the standard UEQ-S short version items:

1. Obstructive—Supportive
2. Complicated—Easy
3. Inefficient—Efficient
4. Confusing—Clear
5. Boring—Exciting
6. Not interesting—Interesting
7. Conventional—Inventive
8. Usual—Leading edge

Appendix C. Prototype Screens




Figure A1. Instagram feed with lock icon to trigger the privacy explanation interface (left), Control option for location sharing appears after the trigger (right).




Figure A2. Participants interact with either Interface A (Contrastive explanation, left) or Interface B (Example-based explanation, right).

References

1. Magoulas, G.S.; Polykalas, S.E. Access to Personal Data is Still Tempting for Mobile Apps Even after the GDPR Implementation. In *Proceedings of the 2023 SEEDA-CECNSM*; IEEE: New York, NY, USA, 2023. <https://doi.org/10.1109/SEEDA-CECNSM61561.2023.10470645>.
2. Li, T.; Xia, T.; Wang, H.; Tu, Z.; Tarkoma, S.; Han, Z.; Hui, P. Smartphone App Usage Analysis: Datasets, Methods, and Applications. *IEEE Commun. Surv. Tutor.* **2022**, *24*, 937–966. <https://doi.org/10.1109/COMST.2022.3163176>.
3. Droste, J.R.C. Development of a Concept for Privacy Explanations and Its Prototypical Evaluation. Master's Thesis, Leibniz University Hannover, Hannover, Germany, 2022.
4. European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council: General Data Protection Regulation. *Off. J. Eur. Union* **2016**, *199*, 1–88.
5. Li, S.; Yang, Z.; Nan, Y.; Yu, S.; Zhu, Q.; Yang, M. Are We Getting Well-informed? An In-depth Study of Runtime Privacy Notice Practice in Mobile Apps. In *Proceedings of the 2024 ACM CCS*; ACM: New York, NY, USA, 2024. <https://doi.org/10.1145/3658644.3670377>.
6. Shen, B.; Wei, L.; Xiang, C.; Wu, Y.; Shen, M.; Zhou, Y.; Jin, X. Can Systems Explain Permissions Better? Understanding Users' Misperceptions under Smartphone Runtime Permission Model. In *Proceedings of the USENIX Security Symposium (USENIX Security 21)*, Online, 11–13 August 2021.
7. Karegar, F.; Pettersson, J.S.; Fischer-Hübner, S. The Dilemma of User Engagement in Privacy Notices. *ACM Trans. Priv. Secur.* **2020**, *23*, 1–38. <https://doi.org/10.1145/3372296>.
8. Deters, H.; Droste, J.; Obaidi, M.; Schneider, K. Exploring the Means to Measure Explainability: Metrics, Heuristics and Questionnaires. *Inf. Softw. Technol.* **2025**, *181*, 107682. <https://doi.org/10.1016/j.infsof.2025.107682>.
9. Pan, S.; Tao, Z.; Hoang, T.; Zhang, D.; Li, T.; Xing, Z.; Xu, X.; Staples, M.; Rakotoarivelo, T.; Lo, D. A NEW HOPE: Contextual Privacy Policies for Mobile Applications and An Approach Toward Automated Generation. In *Proceedings of the 33rd USENIX Security Symposium (USENIX Security 24)*, Philadelphia, PA, USA, 14–16 August 2024.
10. Ebert, N.; Ackermann, K.A.; Scheppeler, B. Bolder is Better: Raising User Awareness through Salient and Concise Privacy Notices. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*; ACM: New York, NY, USA, 2021. <https://doi.org/10.1145/3411764.3445516>.
11. Chazette, L.; Karras, O.; Schneider, K. Do End-Users Want Explanations? Analyzing the Role of Explainability as an Emerging Aspect of Non-Functional Requirements. In *Proceedings of the 2020 IEEE 28th International Requirements Engineering Conference (RE)*; IEEE: New York, NY, USA, 2020. <https://doi.org/10.1109/RE.2019.00032>.
12. Brunotte, W.; Chazette, L.; Korte, K. Can Explanations Support Privacy Awareness? In *Proceedings of the IEEE International Requirements Engineering Workshops (REW '21)*; IEEE: New York, NY, USA, 2021. <https://doi.org/10.1109/REW53955.2021.00032>.
13. Manna, S.; Sett, N. Reconciling Privacy and Explainability in High-Stakes: A Systematic Inquiry. *arXiv* **2025**, arXiv:2412.20798.
14. Lazar, J.; Feng, J.H.; Hochheiser, H. *Research Methods in Human-Computer Interaction*, 2nd ed.; Morgan Kaufmann: Cambridge, MA, USA, 2017.
15. Furini, M.; Mirri, S.; Montangero, M.; Prandi, C. Privacy Perception when Using Smartphone Applications. *Mob. Netw. Appl.* **2020**, *25*, 1055–1061. <https://doi.org/10.1007/s11036-020-01529-z>.
16. Brunotte, W.; Specht, A.; Chazette, L.; Schneider, K. Privacy Explanations—A Means to End-User Trust. *J. Syst. Softw.* **2023**, *195*, 111545. <https://doi.org/10.1016/j.jss.2022.111545>.
17. Chitkara, S.; Gothoskar, N.; Harish, S.; Hong, J.I.; Agarwal, Y. Does this App Really Need My Location? Context-Aware Privacy Management for Smartphones. In *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*; ACM: New York, NY, USA, 2017. <https://doi.org/10.1145/3132029>.
18. Bakopoulou, E.; Shuba, A.; Markopoulou, A. Exposures Exposed: A Measurement and User Study to Assess Mobile Data Privacy in Context. *arXiv* **2022**, arXiv:2008.08973.
19. Prange, S.; Knierim, P.; Knoll, G.; Dietz, F.; Luca, A.D.; Alt, F. I do (not) need that Feature!—Understanding Users' Awareness and Control of Privacy Permissions on Android Smartphones. In *Proceedings of the Twentieth Symposium on Usable Privacy and Security (SOUPS)*, Philadelphia, PA, USA, 12–13 August 2024.
20. Frik, A.; Kim, J.; Sanchez, J.R.; Ma, J. Users' Expectations About and Use of Smartphone Privacy and Security Settings. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*; ACM: New York, NY, USA, 2022. <https://doi.org/10.1145/3491102.3517504>.
21. Brunotte, W.; Chazette, L.; Köhler, L.; Klünder, J.; Schneider, K. What About My Privacy? Helping Users Understand Online Privacy Policies. In *Proceedings of the 30th IEEE International Requirements Engineering Conference (RE)*; IEEE: New York, NY, USA, 2022. <https://doi.org/10.1145/3529320.3529327>.

22. Nguyen, T.T.; Backes, M.; Stock, B. Freely Given Consent? Studying Consent Notice of Third-Party Tracking and Its Violations of GDPR in Android Apps. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20)*; ACM: New York, NY, USA, 2020. <https://doi.org/10.1145/3548606.3560564>.
23. Schaub, F.; Balebako, R.; Durity, A.L.; Cranor, L.F. A Design Space for Effective Privacy Notices. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS)*; ACM: New York, NY, USA, 2015.
24. Windl, M.; Henze, N.; Schmidt, A.; Feger, S.S. Automating Contextual Privacy Policies. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*; ACM: New York, NY, USA, 2022. <https://doi.org/10.1145/3491102.3517688>.
25. McDonald, A.M.; Cranor, L.F. The Cost of Reading Privacy Policies. *I/S J. Law Policy Inf. Soc.* **2008**, *4*, 543.
26. Ortloff, A.M.; Windl, M.; Henze, N.; Schwind, V. Implementation and In Situ Assessment of Contextual Privacy Policies. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*; ACM: New York, NY, USA, 2020. <https://doi.org/10.1145/3357236.3395549>.
27. Chazette, L.; Brunotte, W.; Speith, T. Explainable Software Systems: From Requirements Analysis to System Evaluation. *Requir. Eng.* **2022**, *27*, 457–487. <https://doi.org/10.1007/s00766-022-00393-5>.
28. Chazette, L.; Brunotte, W.; Speith, T. Exploring Explainability: A Definition, a Model, and a Knowledge Catalogue. In *Proceedings of the 2021 IEEE 29th International Requirements Engineering Conference (RE)*; IEEE: New York, NY, USA, 2021. <https://doi.org/10.1109/RE51729.2021.00025>.
29. Chazette, L.; Schneider, K. Explainability as a non-functional requirement: Challenges and recommendations. *Requir. Eng.* **2020**, *25*, 493–514. <https://doi.org/10.1007/s00766-020-00333-1>.
30. van der Waa, J.; Nieuwburg, E.; Cremers, A.; Neerinx, M. Evaluating XAI: A Comparison of Rule-Based and Example-Based Explanations. *Artif. Intell.* **2021**, *291*, 103404. <https://doi.org/10.1016/j.artint.2020.103404>.
31. Alonso, J.M.; Stepin, I.; Pereira-Fariña, M.; Catala, A. A Survey of Contrastive and Counterfactual Explanation Generation Methods for Explainable Artificial Intelligence. *IEEE Access* **2021**, *9*, 11974–12001. <https://doi.org/10.1109/ACCESS.2021.3051315>.
32. Miller, T. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artif. Intell.* **2019**, *267*, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>.
33. Adadi, A.; Berrada, M. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* **2018**, *6*, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>.
34. DIN EN ISO 9241-210:2020-03; Ergonomie der Mensch-System-Interaktion—Teil 210: Menschzentrierte Gestaltung Interaktiver Systeme. Deutsche Fassung. DIN Deutsches Institut für Normung e. V.: Berlin, Germany, 2020. <https://doi.org/10.31030/3104744>.
35. World Medical Association. World Medical Association Declaration of Helsinki: Ethical principles for medical research involving human subjects. *Jama* **2013**, *310*, 2191–2194.
36. Rae, A.; Le, L. *Using Prioritization Matrices to Inform UX Decisions*; Nielsen Norman Group: Dover, DE, USA, 2021.
37. Braun, V.; Clarke, V. Using thematic analysis in psychology. *Qual. Res. Psychol.* **2006**, *3*, 77–101.
38. Patil, S.S.; Goswami, N.; Patil, S.R. Pattern of Social Media Platform Use among Students from Health Sciences University. *J. Datta Meghe Inst. Med. Sci. Univ.* **2024**, *19*, 82–86. https://doi.org/10.4103/jdmimsu.jdmimsu_749_23.
39. Phan, T.H.T.; Pham, N.T.; Hieu, D.M.; Tran, K.D.; Tran, B.Q.; Le, B.K.; Ngan, N.T.K.; Phu, T.N.D. Evaluating Third-Party Involvement in Android Apps: Norms and Anomalies in Usage Patterns. In *Proceedings of the Mobile Web and Intelligent Information Systems: 20th International Conference, MobiWIS 2024, Vienna, Austria, 19–21 August 2024*; Springer: Berlin/Heidelberg, Germany, 2024. https://doi.org/10.1007/978-3-031-68005-2_9.
40. NASA Human Systems Integration Division. *NASA Task Load Index (TLX): Subjective Workload Assessment*; NASA Human Systems Integration Division: Moffett Field, CA, USA, 2023.
41. Schrepp, M.; Hinderks, A.; Thomaschewski, J. Construction of a benchmark for the user experience questionnaire (UEQ). *Int. J. Interact. Multimed. Artif. Intell.* **2017**, *4*, 40–44.
42. Nielsen, J. Quantitative Studies: How Many Users to Test? 2006. Available online: <https://www.nngroup.com/articles/quantitative-studies-how-many-users/> (accessed on 25 June 2026).
43. Ding, Y.; Cao, Y.; Qu, Q.; Duffy, V.G. An Exploratory Study Using Electroencephalography (EEG) to Measure the Smartphone User Experience in the Short Term. *Int. J. Hum.-Interact.* **2020**, *36*, 1008–1021. <https://doi.org/10.1080/10447318.2019.1709330>.
44. Dumas, J.S.; Loring, B.A. *Moderating Usability Tests: Principles and Practices for Interacting*; Elsevier: Amsterdam, The Netherlands, 2008. <https://doi.org/10.1016/B978-0-12-373933-9.X5001-1>.
45. Charness, G.; Gneezy, U.; Kuhn, M.A. Experimental methods: Between-subject and within-subject design. *J. Econ. Behav. Organ.* **2012**, *81*, 1–8. <https://doi.org/10.1016/j.jebo.2011.08.009>.
46. Field, A.; Hole, G. *How to Design and Report Experiments*; SAGE Publications: London, UK, 2003.

47. Bjork, R.; Dunlosky, J.; Kornell, N. Self-Regulated Learning: Beliefs, Techniques, and Illusions. *Annu. Rev. Psychol.* **2012**, *64*, 417–444. <https://doi.org/10.1146/annurev-psych-113011-143823>.
48. Okoyomon, E.; Samarin, N.; Wijesekera, P.; On, A.E.B.; Vallina-Rodriguez, N.; Reyes, I.; Álvaro Feal.; Egelman, S. *On the Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies*; EECS Department, University of California: Berkeley, CA, USA, 2019.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.