UsersFirst in Practice: Evaluating a User-Centric Threat Modeling Taxonomy for Privacy Notice and Choice

Alexandra Xinran Li, Tian Wang, Yu-Ju Yang, Miguel Rivera-Lanas, Debeshi Ghosh, Hana Habib, Lorrie Cranor, and Norman Sadeh

Carnegie Mellon University & University of Illinois Urbana-Champaign

{alexandrali,mriveral,debeshig,htq,lorrie,sadeh}@cmu.edu, {tianw7,yuju2}@illinois.edu

***Abstract*—Privacy regulations impose requirements on data collection and use, including obligations to disclose practices and provide choices free of deceptive patterns, emphasizing user-centric notice and choice delivery. The UsersFirst framework introduces a threat taxonomy to guide organizations in identifying where notices and choices fail to adequately support users. This paper presents an experiment evaluating its effectiveness. Twenty-six participants with privacy expertise analyzed user-centric threats in one of two scenarios, either with or without the taxonomy. Our results show that participants using the taxonomy identified significantly more relevant threats: over twice as many in one scenario and 50% more in the other. While the UsersFirst threat taxonomy helped privacy analysts more effectively identify areas where privacy notices and choice mechanisms fall short, we also identified areas for possible improvements to the taxonomy. Finally, we demonstrate an approach to assessing privacy threat analysis tools that may be useful to other researchers.**

I. Introduction

New privacy and data governance regulations impose increasingly stringent requirements on how data is collected and used [1], [2]. This includes detailed requirements for the disclosure of data practices and the need to offer more comprehensive controls. Penalties for not complying with these requirements have also become significantly steeper [3]. In this context, organizations are looking for guidance to help them systematically identify and mitigate privacy risks, including how to document privacy analysis to help prepare for audits and revisit earlier decisions as regulations or their interpretations evolve [4], [5].

Privacy threat modeling (PTM) frameworks have been introduced to help organizations manage their privacy risk, including LINDDUN [6], [7], the NIST Privacy Risk Management Framework (RMF) [8], MITRE's PANOPTIC framework, and xCOMPASS (Models of Applied Privacy) [9], [10]. Although these frameworks are all useful, they provide limited guidance in identifying usability-related threats that may impact the effectiveness of notice and choice. In contrast, the recently proposed UsersFirst Framework, presented in our prior work,[1] introduces a threat taxonomy to help analysts uncover user-centric threats in privacy notices and choice mechanisms. To our knowledge, it is the first PTM framework specifically focused on user-oriented threats. In prior pilot work with privacy students [11], an early version of UsersFirst helped participants identify such threats in an uncontrolled setting using existing websites. Here, we report on a more systematic study of a more recent version of UsersFirst, where participants are subdivided into a condition that gets the benefit of UsersFirst and one that does not. All participants have professional experience or graduate-level training in privacy.

We conducted a user study evaluating the UsersFirst threat taxonomy in which 26 participants with professional experience or graduate-level education in privacy played the role of privacy analysts tasked with analyzing fictitious product scenarios. The study compares the performance of participants guided to use the user-centric threat taxonomy with participants who completed the same task without the taxonomy. Both sets of participants could take advantage of their prior training and experience in the privacy field and any knowledge they had about other privacy threat modeling frameworks.

Evaluating privacy threat modeling frameworks poses several challenges: recruiting participants with relevant expertise and experience, dedicating sufficient time for comprehensive assessments, and balancing framework coverage with realistic scenarios. Nevertheless, such evaluations are essential to understand how effectively a framework supports analysts. We carefully crafted realistic tasks that could be performed within a 90-minute study by busy privacy professionals, allowing us to assess the effectiveness of the UsersFirst taxonomy.

This paper focuses on the following research questions:

- Does the UsersFirst taxonomy of user-centric threats help people with privacy expertise do a better job identifying threats? Do they identify more relevant threats? Are they more accurate in their identification of relevant threats?
- In what areas does the UsersFirst taxonomy best assist people with privacy expertise, and in what areas is there room for improvement?

[1] https://usersfirst.io

Symposium on Usable Security and Privacy (USEC) 2026
27 February 2026, San Diego, CA, USA
ISBN 978-1-970672-07-7
https://dx.doi.org/10.14722/usec.2026.23012
www.ndss-symposium.org, https://www.usablesecurity.net/USEC/

Our results show that the UsersFirst threat taxonomy helps people with work experience or education in the privacy field identify relevant user-centric notice and choice threats. Participants who performed tasks without the taxonomy missed a substantial number of relevant user-centric notice and choice threats, while those who used the taxonomy identified significantly more. Our study provides the first empirical evidence that UsersFirst increases practitioners' recall without a corresponding loss in precision, while also identifying concrete usability improvements. Finally, we demonstrate an approach to assessing and improving privacy threat analysis tools that may be useful to other researchers.

The remainder of this paper is organized as follows. In Section II, we present background and related work. Section III provides a brief overview of the UsersFirst framework's taxonomy of user-centric threats. Section IV discusses our study methods, including the scenarios, interview design, and data analysis process. Section V presents the results of our study, and Section VI discusses the implications of these results. We present concluding remarks in Section VII.

II. Background and Related Work

In this section, we briefly summarize research on the usability of privacy notices and choices and on existing privacy threat modeling frameworks and their evaluation.

A. Usability of Privacy Notices and Choices

Privacy notices tend to take a long time to read and contain legal and technical jargon that is hard to understand. Studies have shown that they are often vague and difficult to find [12], [13], [14], [15], [16], [17], [18], [19], [20], [21]. An effective short-form privacy notice or privacy "nutrition label" could make it easier for users to obtain privacy information [22], [23], [24], [25] and to overcome the burdens associated with reading long and complex privacy policies [26]. However, not all short-form privacy notices are effective: mobile app privacy nutrition labels in the iOS and Android app stores have been found to also suffer from significant usability problems [27], [28], [29], [30]. Further, no consistent, systematic approach has been applied to identify and document these deficiencies.

Research on privacy choices highlights the need for meaningful, accessible controls that are seamlessly integrated with privacy notices, informative, timely, free of manipulative patterns ("dark patterns" [31]), and aligned with user preferences [32], [33], [34]. Feng et al. identified systemic shortcomings in privacy choice mechanisms, such as a lack of meaningful options, confusing choices, and insufficiently granular controls [32]. Habib et al. analyzed data deletion and opt-out options on websites, finding them inconsistently placed, with missing information and broken links [35]. Several studies have examined cookie banners, identifying design elements that influence user behavior, such as whether users can access cookie-management options on the initial banner [36], [37], [38], [39]. Studies have found that cookie banners tend to push users towards privacy-intrusive choices or make cookie rejection much harder than consent [40], [41], often in violation of regulations [42]. Other studies have examined usability issues related to advertising controls. Garlach et al. found that the AdChoices icon was difficult to find on mobile devices [43], while Im et al. explored design improvements for ad control, focusing on entry points and actionable features [34].

Researchers have also explored the use of privacy profiles as a way of simplifying user choices, especially in contexts where users are otherwise expected to manage an unrealistically large number of privacy choices [44], [45], [46], [47].

Together, these studies demonstrate the challenges in designing privacy notice and choice experiences and the need for actionable guidance for designers of such interfaces. Privacy design guidelines emphasize the need for notice and choice mechanisms that are relevant, actionable, and understandable [33], and recommend following a user-centered design process that includes user testing [48]. Habib and Cranor synthesized published usability evaluations of privacy choice interfaces and proposed a privacy choice evaluation framework that covers seven usability aspects, suggesting study methods to evaluate each aspect [49]. We draw on these methods to assess privacy threat modeling frameworks and their support for designing effective notice and choice mechanisms.

B. Privacy Threat Modeling Frameworks

Privacy threat modeling (PTM) is the process of analyzing a system or service to identify privacy vulnerabilities and ways to prevent or mitigate threats [50]. Solove's privacy taxonomy [51] identifies privacy threats across four categories and provides a foundation for understanding and addressing privacy risks. Nissenbaum's theory of Contextual Integrity [52] can be used to help assess privacy expectations and norms based on social contexts. LINDDUN PRO [53] and PANOPTIC [54] are privacy threat modeling approaches developed to help analysts identify and mitigate a wide range of privacy threats. The NIST Privacy Framework helps organizations manage business-oriented privacy risks [55]. Usman and Zappala [50] proposed a human-centered threat modeling framework to identify security and privacy threats by understanding human needs, perspectives, and experiences. The UsersFirst framework was designed to help practitioners identify threats in the design of privacy notice and choice mechanisms [56]. With the exception of UsersFirst, these frameworks are not specifically designed to surface threats related to user experience with notice and choice interfaces.

Little research has evaluated the use of privacy threat modeling frameworks with privacy practitioners. The most in-depth prior evaluation was Wuyts et al.'s empirical analysis of LINDDUN, which included three descriptive user studies that examined the correctness of the results produced and the ease of use of the framework [57]. The team also evaluated their lightweight toolkit, LINDDUN GO, with students and 10 industry professionals, surveying them on the understandability of the toolkit's threat type cards [6]. However, these studies focused only on how participants used the framework and did not compare participants' performance with LINDDUN to other threat modeling frameworks or without using a

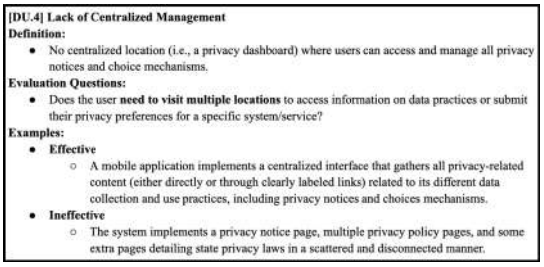

[DU.4] Lack of Centralized Management

Definition:

- No centralized location (i.e., a privacy dashboard) where users can access and manage all privacy notices and choice mechanisms.

Evaluation Questions:

- Does the user **need to visit multiple locations** to access information on data practices or submit their privacy preferences for a specific system/service?

Examples:

- **Effective**
  - A mobile application implements a centralized interface that gathers all privacy-related content (either directly or through clearly labeled links) related to its different data collection and use practices, including privacy notices and choices mechanisms.
- **Ineffective**
  - The system implements a privacy notice page, multiple privacy policy pages, and some extra pages detailing state privacy laws in a scattered and disconnected manner.

Fig. 1. Definition, evaluation questions, and examples for [DU.4] Lack of Centralized Management.

framework. We previously conducted a small study with 14 privacy students comparing threat identification using an early version of the UsersFirst threat taxonomy, LINDDUN PRO's unawareness threat category, and no taxonomy [11]. While we found that the early version of the taxonomy helped privacy students identify threats in an uncontrolled website setting, we also observed that those using LINDDUN PRO performed worse than those using no framework. This motivates our choice to use a no-framework condition rather than an established framework, given that LINDDUN's unawareness threat category is, to our knowledge, the only component of existing frameworks that focuses on user-oriented privacy threats.

In this paper, we introduce a systematic approach to evaluating privacy threat modeling frameworks using an in-depth controlled study with participants with work experience or education in the privacy field. We apply it to an evaluation of the UsersFirst Framework's user-centric threat taxonomy [56].

III. UsersFirst Overview

The UsersFirst Framework is a tool for privacy practitioners involved in the design and evaluation of privacy notices and choice interfaces. The analysis process begins with a *design phase*. During this stage, analysts identify the data practices that should be disclosed to data subjects —including both direct users and non-users — and determine which choices should be made available to them. They also determine the contexts in which different types of users should encounter privacy notices and choices, and specify the "touchpoints" where they will be presented. Finally, they design the notice or choice interfaces for each context. In the *threat analysis* phase, analysts review interfaces using a taxonomy of user-centric threats to inform design revisions intended to mitigate these threats. This phase can be applied to the design or prototype resulting from the design phase, or to a previously designed product. For each notice and choice, the analyst considers each type of user and context and determines whether any of the threats in the user-centric threat taxonomy are present [56].

The UsersFirst taxonomy version 0.9 [56] defines 28 user-oriented threat types in four major threat categories, summarized in Table I. The full version of the threat taxonomy (included in Appendix G) provides threat definitions along with evaluation questions to assist analysts in identifying specific threat types, accompanied by relevant examples. See Figure 1 for an example of how one threat is presented.

TABLE I
UsersFirst threat taxonomy overview. Threats applying only to privacy choices are marked with an asterisk (*). Full taxonomy in Appendix G.

Discovery and Use (DU)
Threats that make it difficult for users to find and use privacy notices and choices.
[DU.1] Nonexistent or Difficult to Locate
[DU.2] Ineffective Timing
[DU.3] Ineffective Channel
[DU.4] Lack of Centralized Management
[DU.5] Decoupled Notice and Choice
[DU.6] Poor Organization
[DU.6.1] Lengthy Text Without Structure
[DU.6.2] Too Much Effort to Access Information
[DU.7] Poor Formatting
[DU.8] Dysfunctional Components
[DU.9] Distracting Visual/Audio Effects
Comprehension (C)
Threats that make notices and choices difficult to understand.
[C.1] Contradictory Statement(s) or Implementation(s)
[C.1.1] Conflicting Statements
[C.1.2] Mismatched Notice vs. Implementation
[C.2] Inconsistent Terminology
[C.3] Difficult to Understand
[C.3.1] Unclear Wording
[C.3.2] Legal/Technical Jargon
[C.3.3] Complex Sentences
[C.4] Consequences Not Explained*
[C.5] Inadequate Feedback*
[C.6] Confusing UI Controls*
Appropriate Choices (AC)
Threats related to types, granularity, or modifiability of choices.
[AC.1] Limited Choice*
[AC.2] Excessive or Redundant Choices*
[AC.3] Inadequate or Excessive Granularity*
[AC.4] Difficult to Modify Previous Choices*
Manipulative Elements (M)
Threats that manipulate users into less privacy-protective actions.
[M.1] Manipulative Statements
[M.2] Visually Manipulative Design*
[M.3] Asymmetric Effort Required*
[M.4] Less Privacy-Protective Defaults*
[M.5] Unexpected Choice Alteration*

IV. Methods

We designed two fictitious scenarios and presented them in the form of storyboards. Each storyboard showed a user persona using a fictional application or device to perform a series of privacy-related tasks involving notices and choices. Our research team collaboratively identified a set of *ground-truth* user-centric privacy threat instances for each scenario using the UserFirst threat taxonomy. We recruited 26 privacy practitioners and individuals with graduate-level training in privacy. We assigned each participant randomly to either a treatment condition, where participants were given the UsersFirst taxonomy and instructions on how to use it, or a control condition without the taxonomy. We randomly assigned participants in each condition to one of the two scenarios and gave them 80 minutes to identify user-centric threats in their assigned scenario. We spent an additional 10 minutes on other parts of the interview. We analyzed the threat instances that participants identified for their relevance to user-centric privacy threats in the context of notices and choices and compared them with the ground-truth instances we identified. This allowed us to assess participant performance and compare outcomes across the treatment and control conditions.

a) Ethical Considerations: This study was approved by the Carnegie Mellon University Institutional Review Board (IRB). All participants received online IRB-approved consent forms before the interview and were asked to confirm their consent during the interview prior to recording. The consent form disclosed that the interview would be conducted via Zoom, a third-party platform with its own privacy policies, transcribed via Otter.ai, and stored online. We assigned each participant a unique ID, which we used to name their recordings, transcripts, and interview documents. We removed personal information and mentions of employers during the transcript cleaning process and stored all materials in an online folder accessible only to research team members.

A. Scenarios

To evaluate how well the UsersFirst taxonomy extends to various technology settings, three researchers reviewed 33 popular mobile applications, IoT devices, and generative AI tools. We selected seven categories of IoT devices that study participants were likely to be familiar with: smart TVs, wearables, voice assistants, smart thermostats, doorbell cameras, smart cameras, and vacuum cleaners. We also focused on mobile apps that integrated technologies such as virtual try-on features and Bluetooth beacons. The researchers reviewed all samples in terms of how personal information was collected and used, examining notices and choices through the lens of the taxonomy's four threat categories. This allowed us to draw inspiration for designing user scenarios that combine realistic design elements with how user-centric threats surface in real-world situations. We decided to focus on mobile apps and IoT devices, for which we found many real-world examples of user-centered privacy threats related to notice and choice.

We designed fictitious mobile applications and IoT device scenarios that included real-world privacy threats that we observed in interfaces for similar applications and devices. To ensure the scenarios reflected realistic rather than artificial threats, we constructed them by combining components from real systems while preserving a threat density comparable to the systems we reviewed. We made minor adjustments: removing certain problematic elements (e.g., overly long privacy policies) and introducing others (e.g., changing the color of toggles) to cover a larger portion of the taxonomy threats. Together, the two scenarios covered 18 of the 28 threat types in the UsersFirst taxonomy. We excluded two categories of threats: those that would not fit naturally into the scenarios, for example, [DU.7] Poorly Formatted Notices and Choices, which would likely be obvious to any analyst but make it difficult to evaluate other threats, and those that would be difficult to evaluate from screenshots, such as [DU.9] Distracting Visual/Audio Effects. We also tried to keep the task manageable within our time constraints and avoid an artificial clustering of threats.

We refer to the mobile application scenario as *AccuFrame* (a mobile app for an eyewear retailer) and to the IoT-device scenario as *Beyond* (a smart TV). Scenarios are described through storyboards where a user persona, Chloe, interacts with different product features and navigates privacy notices and choice mechanisms. Each scenario centers around a privacy-related goal, such as withdrawing previously given consent or enabling a data-control feature.

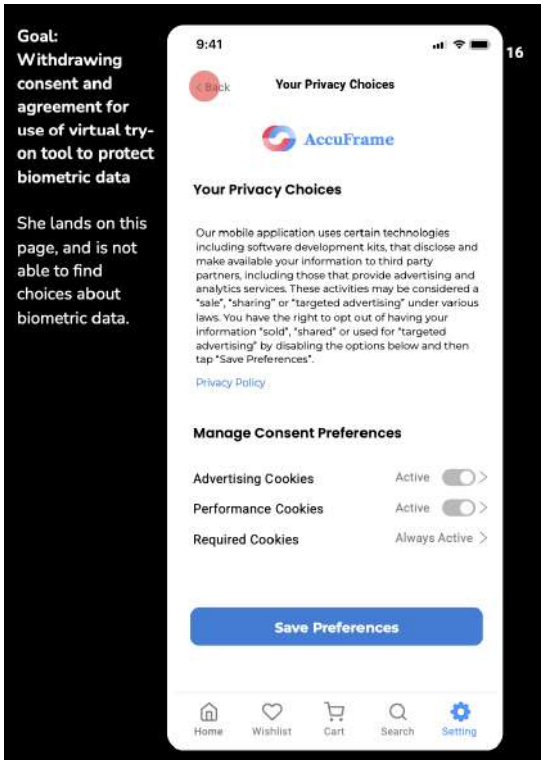

Fig. 2. One page from the AccuFrame storyboard.

To preserve the open-ended nature of the threat-modeling tasks assigned to participants while ensuring that they could complete their tasks in a reasonable amount of time, we limited each scenario to examining a specific set of notices and choices Chloe encountered throughout her user journey. We also asked participants to consider Chloe as the only user for their analysis, rather than considering the needs of other users, such as those who may require accessibility accommodations.

Each privacy notice and choice contained multiple types of threats from the UsersFirst taxonomy. For instance, the cookie management page used in the AccuFrame scenario was designed to include four threat types: DU.6.2, C.3.2, M.2, and M.4 (refer to Table I). As shown in Figure 2, [M.2] Visually Manipulative Design was present in the form of a gray toggle that is shown when an optional cookie is actually active.

During the interview, we provided participants with a 4-5 page introduction, including an overview of their task, definitions of concepts related to privacy notice and choice threats, and a description of the assigned scenario, including the specific list of notices and choices to focus on. The introduction also included instructions for reading the storyboards, followed by a description of Chloe's *user journey* (storyboards) interacting with the product. The user journey includes screenshots, descriptions of Chloe's tasks or goals, relevant thoughts, and explanations of her actions. For exam-

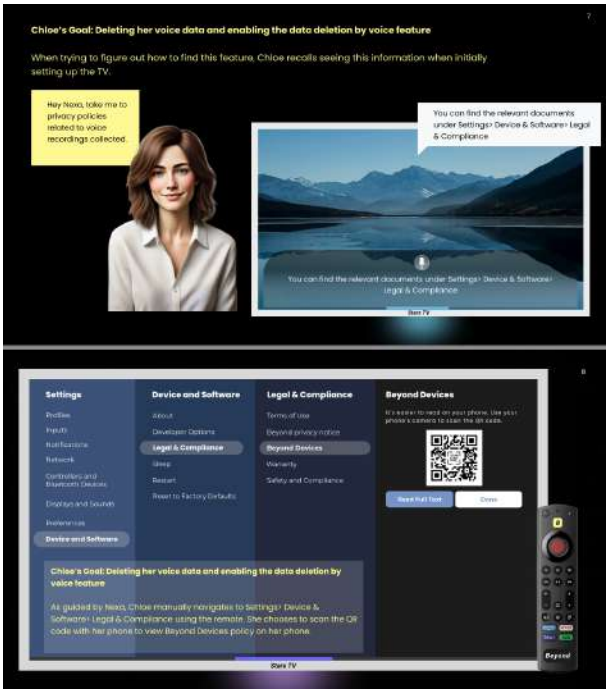

Fig. 3. Two pages from the Beyond storyboard.

ple, Figure 3 features two screens from the Beyond scenario. The top one illustrates Chloe's interaction with the smart TV voice assistant, "Nexa," with her goal displayed at the top, followed by a brief description of her thought process and actions. The smart TV and the response that Chloe receives from "Nexa" are shown on the right side of that page. The second screen highlights Beyond TV's "Settings" page and the navigation path Chloe followed.

When Chloe reaches a privacy policy page, the scenario presents participants with the full text of the long policy document. The scenarios were presented to participants in PDF format to support text searches and copy and paste of text, allowing participants to efficiently document their analysis. The scenario materials and instructions were refined based on feedback from pilot studies. The following provides additional details on the two scenarios.

AccuFrame (Appendix A) is a fictional app for an eyewear retailer and includes a virtual try-on (VTO) feature that allows users to see generated images of themselves trying on glasses. This VTO feature requires users to allow the app to access their device's camera to scan their faces. In this scenario, Chloe uses the VTO feature and is required to consent to the collection and use of her biometric data before the face scan. Later in her journey, she tries to withdraw her consent but fails. Study participants were instructed to focus on the following notices and choices: consent to the collection and use of biometric data, cookie management options, notice about data deletion rights, and the data deletion control page.

Beyond (Appendix B) is a smart TV device that includes an embedded voice assistant "Nexa" that allows users to ask questions and make requests. The smart TV also connects with Nexa's mobile app, which provides users with control over their voice data. In this scenario, Chloe first went through the initial setup of the smart TV. Several months later, she wanted to delete her voice recordings using Nexa but ended up having to navigate to Nexa's mobile app to adjust several settings. Participants were instructed to focus on the following notices and choices: notice about audio data collection and use, notice about deletion of audio data, deletion of voice recording data, management of detected sound history, and control of voice recordings for training data.

B. Recruitment

We recruited participants by email (Appendix C) through privacy professionals' networks. They were required to be at least 18 years old, reside in the US, be able to read long texts in English, and have work experience, graduate education, or executive education in the privacy field. Prospective participants completed a brief screening survey (Appendix D) that included questions about demographics and their privacy expertise. Participants with a suitable privacy background received a Qualtrics survey link with an online consent form, a request for information needed for compensation, and a Calendly link to schedule the study, which took place over Zoom. Upon successful completion of the study, each participant was mailed their chosen compensation (a privacy-themed tote bag or a t-shirt). These options aligned the compensation with the topic of our study and were appealing to our participants.

C. Interview Design

We conducted 90-minute Zoom interviews with 26 participants, with the duration informed by pilot studies conducted beforehand. We required participants to connect to the interview using computers rather than phones or tablets, as the interview tasks involved editing documents online. We also encouraged (but did not require) all participants to bring documents or links related to privacy threat modeling frameworks they were familiar with.

We randomly assigned participants to one of the two scenarios and then further into *with-taxonomy* or *no-taxonomy* treatment groups. This resulted in a total of four condition groups, each with approximately equal numbers of participants: AccuFrame with-taxonomy (*Accu-With*), AccuFrame no-taxonomy (*Accu-No*), Beyond with-taxonomy (*Beyond-With*), and Beyond no-taxonomy (*Beyond-No*). For the with-taxonomy groups, we provided participants with a document containing version 0.9 of the UsersFirst taxonomy [56] (with the name of the framework and other identifying information about the developers of the framework removed) and instructions on how to use it. The no-taxonomy participants were asked to explore freely or use a framework they brought if they preferred to simulate what they would normally do in work settings. We chose the no-taxonomy condition as the control, since to our knowledge, no existing privacy threat modeling framework addresses user-centric threats in the notice and choice context, and our prior work shows that the use of a

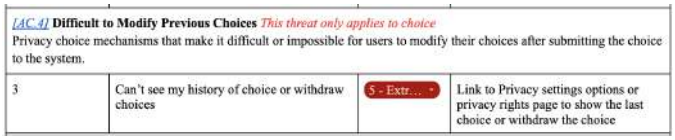

Fig. 4. Sample of a table entry row filled out by a with-taxonomy participant, showing page number, evidence, importance, and design suggestion.

different framework could even hinder participants for this particular type of PTM task [11].

Each interview session had three parts, described below: background questions, threat identification, and threat identification review. Appendix E includes the complete interview scripts for all groups.

Background questions. We began the interview by asking participants to describe their experience in privacy, their familiarity with privacy notice and choice mechanisms, and their prior knowledge of privacy threat modeling.

Threat identification. We briefly introduced participants to their assigned scenario and informed them that their task was to examine the scenario and record user-oriented privacy threats they identified, focusing on the specific list of notices and choices in their assigned scenario. We shared two documents via Zoom: a PDF with the scenario's storyboard (scenario PDF) and a Google Doc to record threats they identified (threat list document).

For the no-taxonomy groups, the threat list document included the list of privacy notices and choices in the assigned scenario with an empty table for each and instructions on what to record in the tables when identifying a threat — including page number, evidence, and importance rating (Appendix F).

Participants in the with-taxonomy groups received similar materials along with a simplified tabular version of the taxonomy with embedded links to a detailed version (Appendix G). We did not tell participants anything about the source of the taxonomy. Instead of giving participants the empty tables provided to the no-taxonomy groups, we included a copy of the simplified taxonomy table for each of the notices and choices and added an empty row under each of the taxonomy threat types (Figure 4). We asked the with-taxonomy participants to record threats they identified under each applicable threat type, adding additional new rows when applicable. For each notice and choice, we provided an "Additional Table" where they could list any threats they felt did not fall under any of the taxonomy threat types, allowing them to avoid placing a threat where it did not fit. The detailed taxonomy and the instructions provided to with-taxonomy participants are included in Appendix G.

To protect participants' privacy and avoid recording any personal information, we asked them to either close existing tabs in their current browser or open a new browser window, then open the scenario PDF and threat list document, and share their screens. We then took 10-15 minutes to walk participants through the two documents, familiarizing them with the user persona's experience using the fictional app or smart TV and UsersFirst's taxonomy (with-taxonomy participants only). We asked participants to identify threats and fill out the threat list table. We also asked participants to explain their thinking and reasoning. Throughout the threat identification process, we asked clarification questions when a participant's explanation was vague or they missed a required field.

Review. In the last part of the interview, we asked participants to comment on their threat-identification experience: how easy or difficult it was to identify user-oriented privacy threats. With-taxonomy participants were asked whether and how the taxonomy had influenced their approach and if they had any suggestions for improving it. No-taxonomy participants were asked to describe their approach to identifying threats and whether there was anything they wished they had during the task that might have helped with their analysis.

D. Data Analysis

1) Ground-Truth Threat Identification: We define a *threat instance* as the specific demonstration of one of the 28 taxonomy threat types in a specific notice or choice — a scenario may include multiple instances of the same threat type. For example, for AccuFrame's Choice 2 (Cookie management options), we identified default cookie acceptance as an instance of the threat [M.4] Less Privacy Protective Defaults. Prior to the interviews, four researchers collaboratively reviewed the scenarios to identify threat instances using the UsersFirst taxonomy. We reached consensus on (1) whether each item constituted a threat instance and (2) which taxonomy threat type it belonged to. Based on this, we developed ground-truth threat instance tables for each scenario. The list of ground-truth threats is shown in Table V in Appendix H-B.

2) Data Preparation and Cleaning: All interviews were recorded, and we used Otter.ai to generate transcripts, timestamps, and screenshots. After each interview, we listened to the recording and edited the transcript to ensure the transcript text and speakers were accurate. Each participant was identified with a unique number according to the group to which they were assigned (e.g., "AN1"). We also removed information that identified participants or their employers from the transcripts. We exported the edited transcripts and stored them in a Google Drive folder only accessible to the research team.

Relevancy coding and threat instance mapping. Two researchers independently coded all participants' threat list documents in two phases.

In the first phase, two coders reviewed all threat instances in the threat lists and collaboratively developed a relevancy codebook to determine whether a privacy threat instance reported by a participant was relevant, e.g., a user-oriented privacy shortcoming of one of the specific notices or choices participants had been instructed to analyze. The codebook also identified "irrelevant" categories to help organize reported threat instances that were not considered relevant (e.g., not privacy related, not focusing on the specified notices and choices, misunderstanding the storyboard).

In the second phase, each coder independently applied the codebook and ground truth to all reported threat instances. When a participant reported multiple threat instances in a

single table entry row, the coders split the entry into multiple reported threat instances. Conversely, when a participant reported the same threat instance under multiple threat types in their table entries, the coders removed these duplicate entries and retained just one of them (deduplication). Following these adjustments, the coders applied the relevancy codebook to determine whether a reported threat instance was relevant (relevancy coding). If a reported threat instance was labeled as irrelevant, we refer to it as an "error." If relevant, the coder documented the corresponding ground-truth threat instance number (e.g., [DU.3] Ineffective Channel for the second choice in the Beyond Scenario). For with-taxonomy participants, the coder also recorded the threat category and type number reported by the participant. We also recorded any relevant threat instances identified by participants that were not included in the ground-truth table, provided that at least two researchers agreed on their relevance. These instances are listed in Table VI in Appendix H-C.

After independently coding each of the participants' reported threat instances, the two coders reconvened and resolved any coding disagreements. If a threat instance identified by a participant was deemed irrelevant, we recorded the type of irrelevance. A third coder reviewed all reported threat instances marked as irrelevant by the first two coders to ensure that no relevant privacy threat instances reported by participants were mistakenly excluded. Agreement was reached between the three coders on the relevance of each reported threat instance.

Placement analysis. For each threat instance identified by with-taxonomy participants that appears in either of the ground-truth threat tables (Table V in Appendix H-B and Table VI in Appendix H-C), we compared the threat type assigned by the participants (i.e., how they categorized the threat instance) with the threat type specified in the ground-truth table. This comparison evaluated whether the participant-assigned category matched our ground truth in both the threat category (e.g., both classified as DU) and threat type (e.g., both classified as DU.2). If a participant documented a threat instance in a notice or choice table but attributed it to the wrong notice or choice (e.g., recording a threat instance related to Choice 1 in the table for Choice 2), we disregarded this difference as long as it was clear that the ground-truth threat instance effectively covered the issue. Additionally, some threat instances can be mapped to multiple taxonomy threat types — e.g., taking a lot of effort to find a privacy control makes it both difficult to find (DU.1) and requires too much effort to access necessary information (DU.6.2). To address the complex nature of the identified privacy threat instances, two coders reviewed the ground-truth tables together and determined the appropriate threat types for each. In the placement analysis, we labeled a threat instance as correctly placed at the threat-category level if the threat category chosen by the participant matches at least one of the ground-truth threat instance's categories. The same rule applies to the analysis at the threat-type level.

Analysis of threat-identification recordings. We used the video recordings of participants working on the threat-identification task to complement our understanding of threat instances when the table entry was vague or if there was a mismatch between the evidence and design suggestions provided by participants. We also reviewed with-taxonomy groups' recordings to observe how they used the taxonomy for threat identification.

Symbol	Definition
T_i	The *total* number of privacy threat instances reported by participant i (including duplicates and errors)
t_i	The number of *unique* privacy threat instances reported by participant i (including errors)
G_i	The *total* number of *ground-truth* threat instances identified by participant i (including duplicates)
g_i	The number of *unique ground-truth* threat instances identified by participant i
G_i^2	The *total* number of *ground-truth* threat instances according to both ground-truth tables (including participant-identified additions) identified by participant i (including duplicates)
$X.g_i$	The number of *unique ground-truth* threat instances in *threat category* X identified by participant i (e.g., DU.g_i for the Discovery and Use category)
$x.g_i$	The number of *unique ground-truth* threat instances in *threat type* x identified by participant i (e.g., DU.1.g_i for DU.1 type)
N_s	The *number* of unique ground-truth threat instances for scenario s (AccuFrame or Beyond)
$X.N_s$	The *number* of unique ground-truth threat instances in *threat category* X for scenario s (e.g., DU.N_A for the number of unique ground-truth threat instances in the DU category for AccuFrame)
$x.N_s$	The *number* of unique ground-truth threat instances in *threat type* x for scenario s (e.g., DU.1.N_A for the number of unique ground-truth threat instances in DU.1 type for AccuFrame)
C_i	The *number* of ground-truth threat instances correctly classified into the *four threat categories*, according to both ground-truth tables (including participant-identified additions), by participant i
c_i	The *number* of ground-truth threat instances correctly classified into the *twenty-eight threat types*, according to both ground-truth tables (including participant-identified additions), by participant i

TABLE II
DEFINITIONS OF NOTATION USED IN CALCULATIONS

3) Review Analysis: Three team members inspected 80% of the data from the review section of the interview and developed two initial codebooks, one for with-taxonomy and one for no-taxonomy. One team member then coded responses from all 26 participants using these codebooks, while two others each independently coded half of the responses. Coders compared and discussed their codes and made adjustments until a consensus was reached on all coding disagreements.

4) Quantitative Analysis: We calculated precision, recall, and correct placement percentages using the definitions below (notation in Table II).

Overall threat identification. To examine whether participants assigned to the with-taxonomy groups (treatment) were able to identify more unique ground-truth instances than those assigned to the no-taxonomy groups (control), we computed the recall (the percentage of total ground-truth threat instances

identified by participant i assigned to scenario s) and precision (the percentage of unique privacy threat instances reported by participant i that were in the ground truth).

$$\text{Recall}_i = \frac{g_i}{\text{N}_\text{s}} \qquad \text{Precision}_i = \frac{g_i}{t_i}$$

We then performed the exact version of the Mann-Whitney U tests, which is appropriate for small sample sizes, on the computed recall and precision by with-taxonomy and no-taxonomy groups for each scenario [58], [59], [60]. We also computed the rank-biserial effect size from the U statistics to determine the strength of the relationship between the variables tested [61].

To determine whether the taxonomy helped participants identify threats they considered important, we also calculated the average importance rating for each taxonomy threat for each scenario. We first determined the average importance rating for each ground-truth threat instance based on participants who identified it within each condition group, and then calculated the overall average for each individual threat type.

To control for false discovery resulting from multiple-testing, we performed post hoc Benjamini-Hochberg adjustment on all p-values [62].

Analyzing threat categories and types. To determine where the taxonomy provided the most assistance, we first computed the number of unique ground-truth instances identified by participant i assigned to scenario s in threat category X or type x divided by the total number of ground-truth instances in that threat category or type, using these formulas:

$$\text{Recall}^4_{s_X\_\text{i}} = \frac{X.g_i}{X.N_s} \qquad \text{Recall}^{28}_{s_x\_\text{i}} = \frac{x.g_i}{x.N_s}$$

We also calculated the average recall per threat category and type for each condition group and made comparisons across with-taxonomy and no-taxonomy groups.

Placement analysis. We calculated the percentage of ground-truth threat instances categorized in that category or type that were correctly placed by participant i according to the four threat categories (X) and 28 threat types (x):

$$\text{Correct Placement}^4_i = \frac{C_i}{G^2_i} \qquad \text{Correct Placement}^{28}_i = \frac{c_i}{G^2_i}$$

For threat types frequently placed incorrectly, we examined where the threat instance should have been placed and looked for potential causes of these mistakes.

V. Results

In this section, we first describe our participants (Section V-A). Next, we present the results of precision and recall analyses along with the threat importance ratings (Sections V-B and V-C). In Section V-D, we examine participants' performance by threat category and individual threat types. Finally, in Section V-E, we describe participants' threat identification experiences, focusing on their approaches and feedback.

A. Participants

We recruited 40 participants to complete our screening survey between August and October 2024; We invited 39 to participate in interviews, and 29 did. After completing all interview sessions, we removed two participants for failing to comprehensively examine all notices and choices, and another for using a tablet to participate in the interview despite the instructions we provided to use a laptop. Our final dataset includes 26 participants, distributed into condition groups as follows: 7 participants in Accu-No (AN), 6 participants in Accu-With (AW), 7 participants in Beyond-No (BN), and 6 participants in Beyond-With (BW).

Participants' ages ranged from 26 to 61 years (M = 36, SD = 9.69). 18 identified as male, 4 as female, and 4 chose not to disclose their gender identity. The most common occupations were Privacy Engineer (10), Software Engineer (5), and Privacy Officer/Attorney (5). Of the 26 participants, 25 had professional experience in privacy-related roles. The remaining participant, currently a PhD student, had previously completed a full-time master's program in privacy engineering. Based on their self-described privacy experience, all participants appeared to be well qualified for the threat identification task. Table IV in Appendix H-A includes more demographic details.

Participants reported encountering interfaces for notice and choice at work (21), in coursework (1), or in daily life (4). Although 23 participants indicated that they understood privacy threat modeling conceptually, only 16 could name specific frameworks, some mentioning more than one: LINDDUN (9), NIST (4), STRIDE (2), MITRE ATT&CK (1), and company-specific frameworks (3). Notably, STRIDE and MITRE ATT&CK are actually primarily used for identifying security threats.

B. Overall Threat Identification

We analyzed each participant's threat tables according to the process detailed in section IV-D1. We computed the average number of unique threat instances reported by participants in each group (Table III) and found that participants in with-taxonomy groups reported an average of 16 threat instances, while that number was slightly lower for participants in no-taxonomy groups (11 participants in the Accu-No group and 14 in Beyond-No group).

The instance mapping process (Section IV-D1) identified 17 unique ground truth threat instances for AccuFrame, spanning 12 taxonomy threat types. For Beyond, 22 unique ground-truth threat instances were identified across 15 taxonomy threat types. Table V in Appendix H-B shows the ground-truth table. We also report a table of threat instances (Table VI in Appendix H-C) that captures relevant threats identified by participants that were not part of our ground truth.

With-taxonomy participants had a significantly higher average recall than no-taxonomy participants for both scenarios ($p = 0.031$, effect size = 0.905 for AccuFrame and $p = 0.047$, effect size = 0.762 for Beyond). Out of the 17 AccuFrame ground-truth threat instances, Accu-With participants successfully identified an average of 10.2, while Accu-No participants identified an average of 5.1. Four out of the six Accu-With participants identified at least half of the ground-truth threat instances, whereas the top-performing participants from Accu-No identified less than half of the ground-truth

Participant Number	Occupation	Used A Framework Brought	Approach	Unique Threat Instances Identified	Unique Ground Truth Threat Instances Identified	Recall	Precision
AW1	Software Engineer	N/A	Scenario	17	9	52.94%	52.94%
AW2	Privacy Engineer	N/A	Taxonomy	19	12	70.59%	63.16%
AW3	Privacy Officer/ Attorney	N/A	Taxonomy	20	10	58.82%	50.00%
AW5	Software Engineer	N/A	Taxonomy	7	7	41.18%	100.00%
AW6	Software Engineer	N/A	Taxonomy	22	16	94.12%	72.73%
AW7	Privacy Product Manager	N/A	Taxonomy	9	7	41.18%	77.78%
AW Average				**15.7**	**10.2**	**59.80%**	**69.43%**
AN1	Software Engineer	No		8	5	29.41%	62.50%
AN2	PhD Student/ Trained as Privacy Engineer	No		13	3	17.56%	23.08%
AN3	Privacy Engineer	No		7	5	29.41%	71.43%
AN4	Privacy Engineer	No		12	4	23.53%	33.33%
AN5	Privacy Engineer	Yes		12	6	35.29%	50.00%
AN6	Past Privacy Officer/ Attorney	No		10	5	29.41%	50.00%
AN7	Privacy Engineer	No		15	8	47.06%	53.33%
AN Average				**11.0**	**5.1**	**30.25%**	**49.1%**
BW1	Software Engineer	N/A	Scenario	20	14	63.34%	70.00%
BW3	Privacy Officer/ Attorney	N/A	Taxonomy	21	11	50.00%	52.38%
BW4	Privacy Product Manager	N/A	Scenario	12	9	40.91%	75.00%
BW5	Privacy Engineer	N/A	Taxonomy	15	11	50.00%	73.33%
BW7	Security Engineer	N/A	Taxonomy	14	7	31.82%	50.00%
BW8	Privacy Engineer	N/A	Scenario	13	9	40.91%	69.23%
BW Average				**15.8**	**10.2**	**46.21%**	**64.99%**
BN1	Privacy Engineer	No		14	9	40.91%	64.29%
BN2	Software Engineer	Yes		12	6	27.27%	50.00%
BN3	Privacy Officer/ Attorney	No		10	9	40.91%	90.00%
BN4	Privacy Engineer	No		10	5	22.73%	50.00%
BN5	Privacy Officer/ Attorney	Yes		11	6	27.27%	54.55%
BN6	Privacy Engineer	No		22	8	36.36%	36.36%
BN7	Privacy Officer/ Attorney	No		18	6	27.27%	33.33%
BN Average				**13.9**	**7.0**	**31.82%**	**54.08%**

TABLE III

PARTICIPANTS' OCCUPATION, THEIR USE OF FRAMEWORK DURING THE INTERVIEW, NUMBER OF UNIQUE THREAT INSTANCES IDENTIFIED, UNIQUE GROUND TRUTH THREAT INSTANCES IDENTIFIED, RECALL, AND PRECISION AS WELL AS AVERAGE FOR EACH CONDITION GROUP. *Scenario* REFERS TO SCENARIO-CENTRIC APPROACHES AND *Taxonomy* REFERS TO TAXONOMY-CENTRIC APPROACHES. ACCU-WITH PARTICIPANTS HAVE THE PREFIX AW, ACCU-NO PARTICIPANTS HAVE THE PREFIX AN, BEYOND WITH PARTICIPANTS HAVE THE PREFIX BW, AND BEYOND-NO PARTICIPANTS HAVE THE PREFIX BN.

threat instances. The performance difference between the with-taxonomy and no-taxonomy groups was also evident in the Beyond scenario: Beyond-With participants identified an average of 10.2 out of 22 ground-truth threat instances, compared to an average of 7 identified by Beyond-No participants. All Beyond-With participants identified between 9 and 14 ground-truth threat instances. In contrast, only two of the seven Beyond-No participants identified a number of ground-truth threat instances within the same range, while the remaining five participants identified fewer than 9 ground-truth threat instances. Table III includes the overall threat identification results, participants' occupation, and whether they used the framework they brought.

The with-taxonomy groups achieved an average precision (percentage of reported threat instances covered by ground-truth threat instances) of 69.4% (Accu-With) and 65.0% (Beyond-With), while the no-taxonomy groups have a lower average (49.1% for Accu-No and 54.1% for Beyond-No). The difference in precision between the with-taxonomy and no-taxonomy groups is not statistically significant in either scenario.

C. Threat Importance Ratings

We computed the average importance rating of a threat type for all ground-truth threat types to assess whether participants who reported threats thought they were important. We found that participants rated almost all threat types as at least moderately important for both scenarios (see Table VII in

Appendix H-D). This suggests that when participants reported threats, they usually found them fairly important. However, we note that participants only provided importance ratings for instances of threats they reported; participants may not have bothered to report some threats that they viewed as unimportant. AccuFrame participants labeled all ground-truth threat types with an average importance rating of at least 3 (moderately important). Accu-With participants rated 3 out of 12 threat types as at least 4 ("very important"), while Accu-No participants rated 5 out of 12 threat types as very important. Three threats were labelled as very important by both groups — [C.1.1] Conflicting Statement(s), [AC.1] Limited Choice, and [AC.4] Difficult to Modify Previous Choices — indicating that certain threats are viewed as important regardless of whether participants used the taxonomy.

For Beyond, all threat types besides [AC.3] Inadequate or Excessive Granularity have an average importance rating of at least 3. Similar to the AccuFrame participants, Beyond participants also rated [C.1.1] Conflicting Statement(s) as very important, and additionally identified [M.3] Asymmetric Effort Required for Different Privacy Protection Levels as very important. Both Beyond-With and Beyond-No groups gave [AC.3] Inadequate or Excessive Granularity an average importance rating of 2. This may be because the threat is not linked to a specific feature,and is instead more subjective and dependent on individual user preferences. We found no clear relationship between average importance and the percentage of participants who successfully identified the threat.

D. Threat Categories and Types

For all four threat categories in AccuFrame, the UsersFirst threat taxonomy helped with-taxonomy participants identify at least some threat instances (see Table VIII and Table IX in Appendix H-E). By contrast, every Accu-No participant missed at least one category, with three participants missing two or more.

Four out of the six Beyond-With participants were able to identify at least some issues from each of the four categories, while four out of the seven Beyond-No participants missed some of the categories. We observed that five Accu-No participants, two Beyond-With participants, and three Beyond-No participants failed to identify any threats in the Appropriate Choice (AC) category, likely due to the fact that only two ground-truth threat instances existed in that category across both scenarios. Figure 5 shows the trends in recall by threat category and condition group, suggesting that some categories of threats were easier for participants to identify both with and without the taxonomy.

We conducted similar recall analyses per ground-truth threat types for both scenarios. For most threat types, participants from with-taxonomy groups were more likely to identify relevant threat instances than no-taxonomy participants (Table X in Appendix H-F).

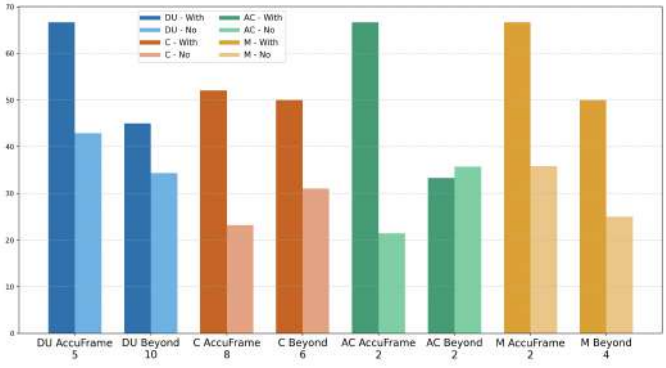

Fig. 5. Recall comparison between with-taxonomy (treatment) and no-taxonomy (control) participants across scenarios and threat categories. Numbers below each threat category indicate the total ground-truth threats for that category in the corresponding scenario.

E. Threat Identification Experience

In this section, we present participants' threat identification experiences (see codebook in Appendix H-H).

1) Approach for Threat Identification: Half of the participants from **no-taxonomy** groups (7 out of 14) approached the threat identification process by taking on specific roles, including user, privacy analyst, legal team, or business. As BN7 explained, "My process is reading everything, toggling between being from a business perspective, a regulator perspective, and from a consumer perspective." Many no-taxonomy participants (9 out of 14) took the approach of considering whether the components presented in the scenario aligned with their perceptions of best and worst practices. For example, they looked for contradictions between the privacy policy and implementation, use of jargon, existence of manipulative statements, potential misuse of data, and privacy controls. For instance, AN6 noted, "The first consent page, it is full of legalese. I don't understand anything that's being said."

Five of the twelve participants in the no-taxonomy condition brought an existing framework with them. However, only one participant from Accu-No and two from Beyond-No actually used their frameworks during the study (see Table III). The three who used their frameworks —LINDDUN, STRIDE+OECD principles, and NIST 800-53 — did not strictly adhere to them, instead relying on their own expertise, similar to other no-taxonomy participants. We did not observe their recall or precision to be outliers within their groups.

For **with-taxonomy** participants, we analyzed and categorized the approach they took during the threat identification process by reviewing their recordings. We observed 7 out of 12 with-taxonomy participants taking a "taxonomy-centric approach," using the UsersFirst threat taxonomy as a checklist. Participants following this approach reviewed each threat type systematically to determine whether they had encountered an example of it while examining the scenario. We observed 5 out of 12 with-taxonomy participants taking a "scenario-centric approach," primarily focusing on the scenario PDF to

identify instances of threats first, and then matching them to a corresponding taxonomy threat type.

7 out of 12 with-taxonomy participants said that the taxonomy helped them identify threats that they might not have otherwise considered, with five of them describing the taxonomy as easy and intuitive to use. AW6 provided an example, stating that "unexpected choice alteration is something that I wouldn't have immediately thought of." Three participants, however, described the taxonomy as a bit heavy. For instance, BW1 stated, "There are just too many threats at a time, like twentyish." Two participants noted that since they were already aware of the threats listed in the taxonomy, their approach to the threat identification process did not change much.

2) Feedback on Threat Identification Process: **No-taxonomy** participants referenced their prior experiences with threat identification, with two participants stating that they had performed similar analyses as privacy practitioners. As AN6 noted, "I lead a team that does this." They also shared their perception of the exercise's ease or difficulty. 7 out of 14 no-taxonomy participants considered the process easy or straightforward, noting no particular challenge throughout the process. Six other participants noted a mix of easy and difficult aspects, while one struggled with the exercise and criticized the LINDDUN framework he brought. Some participants complained that the scenarios themselves were too long, complicated, vague, or insufficient to convey user-oriented threats accurately.

Six participants from no-taxonomy groups suggested the need for a framework or a more useful framework to guide the identification process. AN1 noted, "Maybe if I had, like, some, like threat modeling framework...let's say, have a predefined checklist." BN3 also highlighted the value of having predefined lists, stating, "If I kind of had an outline, here are some potential threats, you know, find throughout this."

For the **with-taxonomy** groups, five participants praised the taxonomy as a helpful checklist or guide. As AW2 noted, "I think it was very useful for ... giving a good list of things to look out for." Likewise, BW3 described it as providing "good prompts ... to use as kind of like a checklist afterward to see if they missed anything from their issue spotting." Other positive feedback on the taxonomy included being detailed and well-explained (4) and being more useful than other frameworks (3). For example, AW7 stated that the taxonomy is "more applicable and efficient when it comes to use compared to other existing frameworks (LINDDUN)." and commented that LINDDUN was focused more on database-related concerns, and UsersFirst is more user-centric and better suited for the threat-identification task.

However, one of the major criticisms was the repetitive format of the tables we provided for recording threats, which four participants suggested merging into a single table. Two participants also suggested improving the taxonomy by broadening its scope to include business, legal, and design perspectives. Two participants complained about overlapping threats, and one felt that the threats were detached from actual harms. Some also raised issues related to the complexity of the interview setup and the effort required to familiarize themselves with the taxonomy, complaining that it was difficult to digest the materials provided in the limited time.

VI. Discussion

In this section, we discuss study limitations, the observed advantages of using the taxonomy for threat identification, issues with determining the ground truth, and suggestions for improving the taxonomy.

A. Limitations

Prior threat-modeling frameworks have not been rigorously evaluated in experimental user studies, likely due to the complexity of conducting such evaluations. An ideal evaluation would involve practitioners using a framework multiple times in real work settings, but this would require industry cooperation and may not be feasible to run as a controlled experiment. We designed our study to maximize ecological validity within the constraints we had. However, due to time limitations, participants had to perform assigned tasks while simultaneously familiarizing themselves with the taxonomy and scenarios. Thus, they may not have had adequate time for the assigned task, and we were unable to measure any learning effects that might occur over time.

For with-taxonomy participants especially, the 90-minute time limit posed a challenge as they had to quickly familiarize themselves with an information-dense taxonomy and apply it systematically within an unfamiliar scenario. Participants said that they might have been able to perform a more thorough analysis if they had reviewed the taxonomy prior to the interview or had additional time. The time limit appeared to have a larger impact on Beyond than on AccuFrame, as the Beyond scenario involved more devices, notices, and choices. Researchers conducting similar studies should allocate more time to threat identification or use simpler scenarios.

We had participants observe a persona's actions rather than explore a real system themselves. As a result, they may have overlooked threats such as [DU.6.2] Too Much Effort to Access Necessary Information because they did not personally experience the effort involved. Chloe's eventual success in locating controls, even with difficulty, may have also made the process seem less problematic. This limitation likely affected all participants similarly, regardless of taxonomy use.

We did not embed all 28 UsersFirst threat types in the scenarios in order to preserve realism and avoid making the scenarios appear contrived.

Our participants had professional or educational backgrounds in privacy but were not specialized threat-modeling experts; thus, our findings reflect how the taxonomy supports general privacy professionals rather than seasoned threat-modeling practitioners.

B. Support for practitioners with varying experience

Participants in the with-taxonomy groups systematically examined user-oriented privacy threats by taking either a taxonomy- or scenario- centric approach. We observed that

participants with privacy threat modeling work experience tended to take a scenario-centric approach and were able to quickly make the connections between a threat instance they found in the scenario and the corresponding taxonomy threat types. They became familiar with the threat identification process faster and usually used the taxonomy as a checklist at the end to identify any threat types they missed. Participants with less experience used the taxonomy from the start in a time-consuming process. As these participants grew more familiar with the taxonomy, they were able to work faster.

C. Is There Really a Ground Truth?

We established ground-truth threat instances to quantify whether the taxonomy added value beyond participants' usual approaches. However, analysts may disagree on whether some ground-truth instances represent threats. We erred on the side of counting identified instances as threats if they were relevant. However, whether or not some of these instances are problematic is a matter of opinion and may require context not provided by the scenarios.

During coding, we considered threats surfaced by participants that were not in our original list. While many of these were relevant, they were typically identified by only one or two participants, likely suggesting that most participants did not view them as important enough to report. While those who did report them often rated them as important, we cannot know how non-reporting participants would have rated them. Our decision to exclude such low-consensus threats from the ground truth reduced precision rates across all conditions, but provided a more consistent baseline for comparison.

We expect that practitioners deploying the taxonomy at work would use it to identify potential threats, and would debate with colleagues, consult legal guidance, or use their judgment to decide which threats should be remediated.

D. Areas for Taxonomy Improvement

Our study provided rich data about how participants used UsersFirst's threat taxonomy, as we had the opportunity to observe participants in real-time and ask them to reflect on their experiences. This allows us to propose refinements to improve the taxonomy. We observed many cases where participants correctly identified something as an instance of a privacy threat but did not correctly determine which threat type it was. Often, their reported threat label suggested a slight misreading of the threat definition or confusion between related threats. This suggests that the UsersFirst taxonomy may be overly granular, and several related threat types should be combined, especially those with similar mitigation.

Additionally, certain threat types may benefit from more precise naming, clearer definitions, and more explicit guidance on their application. For instance, the UsersFirst taxonomy defines certain threats as applicable only to choices. However, participants associated [C.4], Consequences Not Adequately Explained — intended as a choice-specific threat — with unclear consequences described in notices (e.g., "Unclear how data has been handled"). This suggests a need to clarify the scope and context in which each threat type applies.

Similarly, varied interpretations of "privacy notice" and the definition of [DU.1] Nonexistent or Difficult to Locate led participants to misplace several threat instances. According to the taxonomy, [DU.1] refers to the notice itself (e.g., a privacy policy) being missing or hard to find, not missing disclosures within the policy. However, many participants treated absent details about data practices (e.g., storage, provider contact information) as DU.1 threats. To clarify this confusion, we recommend redefining DU.1 by separating the "difficult to locate" part from "non-existent," updating the latter to encompass any missing content/disclosure regarding data practices in the privacy notice and choice, and combining the former with [DU.6.2] Too Much Effort to Access Necessary Information.

Participants also frequently placed a threat instance incorrectly under the Manipulative Elements category (M), suggesting that threat types under this category are confusing. For instance, [M.2] Visually Manipulative Design overlaps with [C.6] Confusing Buttons/Toggles/Checkboxes, as both involve components that can steer users toward less privacy-protective actions. Future taxonomy revisions could distinguish the unique aspects of each type, for example, emphasizing C.6's focus on unclear or misleading action pathways, while reserving M.2 for manipulative strategies.

Since the completion of this study, we released an improved version of the taxonomy incorporating these lessons [2].

VII. Conclusions

We evaluated UsersFirst's user-centric privacy threat taxonomy by conducting a study with 26 individuals with privacy work experience or education. Participants were tasked with analyzing privacy threats in one of two scenarios, either with or without using the taxonomy. While participants reported a similar number of threats regardless of whether they used the taxonomy, participants who used the taxonomy were able to identify significantly more ground-truth threat instances in both scenarios (AccuFrame: 59.8% vs. 30.3% $p = 0.031$; Beyond 46.2% vs 31.8% $p = 0.047$), and did so without any loss in precision. Participants rated all but one ground-truth threat instance as at least moderately important. Taken together, these results suggest that the taxonomy helped participants systematically identify user-oriented privacy threats that were both relevant and important. Our work is among the first to present an evaluation method through observation of privacy practitioners carrying out threat identification tasks. It demonstrates that this approach can surface areas for improvement and offers insight into designing scenarios for testing and refining such frameworks, while highlighting the inherent difficulty of threat identification. Our evaluation methods, based on the creation of targeted scenarios and threat identification tasks, could also be adapted to the evaluation of other future user-centric privacy threat taxonomies.

[2] https://www.usersfirst.io/user-oriented-privacy-threats

Acknowledgment

The researchers gratefully acknowledge the support of the Digital Transformation and Innovation Center at Carnegie Mellon University, sponsored by PwC. Work on the UsersFirst framework has also been supported in part by a grant from Meta.

References

[1] E. Commission, *EU Data Protection Rules*, April 2016, https://commission.europa.eu/law/law-topic/data-protection/eu-data-protection-rules_en.

[2] C. L. Information, *California Consumer Privacy Act (CCPA)*, 2018, https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5.

[3] J. Wolff and N. Atallah, "Early gdpr penalties: Analysis of implementation and fines through may 2020," *Journal of Information Policy*, vol. 11, pp. 63–103, 2021.

[4] R. S. Ross, "Risk management framework for information systems and organizations: A system life cycle approach for security and privacy," National Institute of Standards and Technology, Tech. Rep., 2024.

[5] R. Galvez and S. Gurses, "The odyssey: Modeling privacy threats in a brave new world," in *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2018, pp. 87–94.

[6] K. Wuyts, L. Sion, and W. Joosen, "Linddun go: A lightweight approach to privacy threat modeling," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 302–309.

[7] K. Wuyts and W. Joosen, "Linddun privacy threat modeling: a tutorial," *CW Reports*, 2015.

[8] NIST, "Nist risk management framework (rmf) small enterprise quick start guide: A comprehensive, flexible, risk-based approach to managing information security and privacy risk," 2024.

[9] S. Katcher, B. Ballard, C. Bloom, K. Isaacson, J. McEwen, S. Shapiro, S. Slotter, M. Paes, and R. Xu, "The panoptic™ privacy threat model," 2024.

[10] J. Dev, B. Rashidi, and V. Garg, "Models of applied privacy (map): A persona based approach to threat modeling," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–15.

[11] X. A. Li, Y.-J. Yang, Y. Maurya, T. Wang, H. Habib, N. Sadeh, and L. F. Cranor, "Design and evaluation of the usersfirst privacy notice and choice threat analysis taxonomy," in *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024) Poster Session*, 2024. [Online]. Available: https://www.usenix.org/conference/soups2024/presentation/li-poster

[12] J. R. Reidenberg, T. Breaux, L. F. Cranor, B. French, A. Grannis, J. T. Graves, F. Liu, A. McDonald, T. B. Norton, R. Ramanath, C. Russell, N. Sadeh, and F. Schaub, "Disagreeable privacy policies: Mismatches between meaning and users' understanding," *Berkeley Tech. LJ*, vol. 30, p. 39, 2015.

[13] A. J. Perez, S. Zeadally, and J. Cochran, "A review and an empirical analysis of privacy policy and notices for consumer internet of things," *Security and Privacy*, vol. 1, no. 3, p. e15, 2018.

[14] L. F. Cranor, "Necessary but not sufficient: Standardized mechanisms for privacy notice and choice," *J. on Telecomm. & High Tech. L.*, vol. 10, p. 273, 2012.

[15] B. Fabian, T. Ermakova, and T. Lentz, "Large-scale readability analysis of privacy policies," in *Proceedings of the international conference on web intelligence*, 2017, pp. 18–25.

[16] P. G. Leon, J. Cranshaw, L. F. Cranor, J. Graves, M. Hastak, B. Ur, and G. Xu, "What do online behavioral advertising privacy disclosures communicate to users?" in *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*, 2012, pp. 19–30.

[17] F. H. Cate, "The limits of notice and choice," *IEEE Security & Privacy*, vol. 8, no. 2, pp. 59–62, 2010.

[18] C. Jensen and C. Potts, "Privacy policies as decision-making tools: an evaluation of online privacy notices," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 471–478. [Online]. Available: https://doi.org/10.1145/985692.985752

[19] P. M. Schwartz and D. Solove, "Notice & choice," in *The Second NPLAN/BMSG Meeting on Digital Media and Marketing to Children*, vol. 7, 2009, pp. 1–6.

[20] R. Chen, F. Fang, T. Norton, A. M. McDonald, and N. Sadeh, "Fighting the fog: Evaluating the clarity of privacy disclosures in the age of ccpa," in *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 2021, pp. 73–102.

[21] A. Ravichander, I. Yang, R. Chen, S. Wilson, T. Norton, and N. Sadeh, "Incorporating taxonomic reasoning and regulatory knowledge into automated privacy question answering," in *Web Information Systems Engineering – WISE 2024*. Singapore: Springer Nature Singapore, 2025, pp. 444–460.

[22] M. M. Ali, D. G. Balash, M. Kodwani, C. Kanich, and A. J. Aviv, "Honesty is the best policy: On the accuracy of apple privacy labels compared to apps' privacy policies," *arXiv preprint arXiv:2306.17063*, 2023.

[23] Y. Meier, J. Schäwel, and N. C. Krämer, "The shorter the better? effects of privacy policy length on online privacy decision-making," *Media and Communication*, vol. 8, no. 2, pp. 291–301, 2020.

[24] R. Khandelwal, A. Nayak, P. Chung, K. Fawaz, A. Bianchi, Z. B. Celik, Y. Yarom, X. S. Shen, Z. Fang, S. Zhang *et al.*, "Unpacking privacy labels: A measurement and developer perspective on google's data safety section," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 2831–2848.

[25] V. Jain, S. Ghanavati, S. T. Peddinti, and C. McMillan, "Towards fine-grained localization of privacy behaviors," in *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2023, pp. 258–277.

[26] D. G. Balash, M. M. Ali, C. Kanich, and A. J. Aviv, "" i would not install an app with this label": Privacy label impact on risk perception and willingness to install {iOS} apps," in *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, 2024, pp. 413–432.

[27] S. Zhang, Y. Feng, Y. Yao, L. F. Cranor, and N. Sadeh, "How usable are ios app privacy labels?" *Proceedings on Privacy Enhancing Technologies*, 2022.

[28] S. Zhang, L. Klucinec, K. Norton, N. Sadeh, and L. F. Cranor, "Exploring Expandable-Grid designs to make iOS app privacy labels more usable," in *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 139–157. [Online]. Available: https://www.usenix.org/conference/soups2024/presentation/zhang

[29] Y. Lin, J. Juneja, E. Birrell, and L. F. Cranor, "Data safety vs. app privacy: Comparing the usability of android and ios privacy labels," *arXiv preprint arXiv:2312.03918*, 2023.

[30] Y. Li, D. Chen, T. Li, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding ios privacy nutrition labels: An exploratory large-scale analysis of app store data," in *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 2022, pp. 1–7.

[31] C. M. Gray, C. T. Santos, N. Bielova, and T. Mildner, "An ontology of dark patterns knowledge: Foundations, definitions, and a pathway for shared knowledge-building," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, ser. CHI '24. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: https://doi.org/10.1145/3613904.3642436

[32] Y. Feng, Y. Yao, and N. Sadeh, "A design space for privacy choices: Towards meaningful privacy control in the internet of things," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: https://doi.org/10.1145/3411764.3445148

[33] F. Schaub, R. Balebako, and L. F. Cranor, "Designing effective privacy notices and controls," *IEEE Internet Computing*, vol. 21, no. 3, pp. 70–77, 2017.

[34] J. Im, R. Wang, W. Lyu, N. Cook, H. Habib, L. F. Cranor, N. Banovic, and F. Schaub, "Less is not more: Improving findability and actionability of privacy controls for online behavioral advertising," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3544548.3580773

[35] H. Habib, Y. Zou, A. Jannu, N. Sridhar, C. Swoopes, A. Acquisti, L. F. Cranor, N. Sadeh, and F. Schaub, "An empirical analysis of data deletion and {Opt-Out} choices on 150 websites," in *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, 2019, pp. 387–406.

[36] H. Habib, M. Li, E. Young, and L. Cranor, ""okay, whatever": An evaluation of cookie consent interfaces," in *Proceedings of the 2022 CHI conference on human factors in computing systems*, 2022, pp. 1–27.

[37] E. R. Bouma-Sims, M. Li, Y. Lin, A. Sakura-Lemessy, A. Nisenoff, E. Young, E. Birrell, L. F. Cranor, and H. Habib, "A us-uk usability evaluation of consent management platform cookie consent interface design on desktop and mobile," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3544548.3580725

[38] C. Utz, M. Degeling, S. Fahl, F. Schaub, and T. Holz, "(un) informed consent: Studying gdpr consent notices in the field," in *Proceedings of the 2019 acm sigsac conference on computer and communications security*, 2019, pp. 973–990.

[39] M. Nouwens, I. Liccardi, M. Veale, D. Karger, and L. Kagal, "Dark patterns after the gdpr: Scraping consent pop-ups and demonstrating their influence," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–13.

[40] G. Kampanos and S. F. Shahandashti, "Accept all: The landscape of cookie banners in greece and the uk," in *IFIP International Conference on ICT Systems Security and Privacy Protection*. Springer, 2021, pp. 213–227.

[41] C. Santos, A. Rossi, L. Sanchez Chamorro, K. Bongard-Blanchy, and R. Abu-Salma, "Cookie banners, what's the purpose? analyzing cookie banner text through a legal lens," in *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 2021, pp. 187–194.

[42] C. Matte, N. Bielova, and C. Santos, "Do cookie banners respect my choice?: Measuring legal compliance of banners from iab europe's transparency and consent framework," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 791–809.

[43] S. Garlach and D. D. Suthers, "I'm supposed to see that?'adchoices usability in the mobile environment," 2018.

[44] B. Liu, J. Lin, and N. Sadeh, "Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help?" in *Proceedings of the 23rd international conference on World Wide Web*, 2014, pp. 201–212.

[45] J. Lin, B. Liu, N. Sadeh, and J. I. Hong, "Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings," in *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, 2014, pp. 199–212.

[46] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow my recommendations: A personalized privacy assistant for mobile app permissions," in *Twelfth symposium on usable privacy and security (SOUPS 2016)*, 2016, pp. 27–41.

[47] S. Zhang, Y. Feng, A. Das, L. Bauer, L. F. Cranor, and N. Sadeh, "Understanding people's privacy attitudes towards video analytics technologies," *Proceedings of the FTC PrivacyCon*, pp. 1–18, 2020.

[48] F. Schaub and L. F. Cranor, "Usable and useful privacy interfaces," *An introduction to privacy for technology professionals*, pp. 176–299, 2020.

[49] H. Habib and L. F. Cranor, "Evaluating the usability of privacy choice mechanisms," in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022, pp. 273–289.

[50] W. Usman and D. Zappala, " SoK: A Framework and Guide for Human-Centered Threat Modeling in Security and Privacy Research ," in *2025 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2025, pp. 33–33. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/SP61157.2025.00033

[51] D. J. Solove, *Understanding privacy*. Harvard university press, 2010.

[52] H. Nissenbaum, "Privacy as contextual integrity," *Wash. L. Rev.*, vol. 79, p. 119, 2004.

[53] L. Sion and W. Joosen, "Linddun pro privacy threat modeling tutorial," Tech. Rep., 2023. [Online]. Available: http://www.linddun.org/pro

[54] S. Katcher, B. Ballard, C. Bloom, K. Isaacson, J. McEwen, S. Shapiro, S. Slotter, M. Paes, and R. Xu, "The panoptic™ privacy threat model," in *Twentieth Symposium on Usable Privacy and Security (SOUPS)*, 2024.

[55] K. R. Boeckl and N. B. Lefkovitz, "Nist privacy framework: A tool for improving privacy through enterprise risk management, version 1.0," 2020.

[56] N. Sadeh, L. Cranor, H. Habib, T. Wang, A. X. Li, G. Gopi, S. Ko, Y. Maurya, Y. Qiu, M. Rivera-Lanas, and Z. Song, *UsersFirst Threat Taxonomy, Version 0.9*, May 2024, https://www.usersfirst.io/frameworks/framework-0-9.

[57] K. Wuyts, R. Scandariato, and W. Joosen, "Empirical evaluation of a privacy-focused threat modeling methodology," *Journal of Systems and Software*, vol. 96, pp. 122–138, 2014.

[58] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The annals of mathematical statistics*, pp. 50–60, 1947.

[59] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in statistics: Methodology and distribution*. Springer, 1992, pp. 196–202.

[60] N. Nachar *et al.*, "The mann-whitney u: A test for assessing whether two independent samples come from the same distribution," *Tutorials in quantitative Methods for Psychology*, vol. 4, no. 1, pp. 13–20, 2008.

[61] M. Tomczak and E. Tomczak, "The need to report effect size estimates revisited. an overview of some recommended measures of effect size," 2014.

[62] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal statistical society: series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.

APPENDIX A
STORYBOARD - ACCUFRAME

i

Background Info

Privacy Notice
A privacy notice is a presentation of terms, sometimes but not exclusively in the form of text in a privacy policy or terms of use agreement, intended to inform users about the data practices of a system and what rights, if any, a user of the system might be able to exercise.

Privacy Choice
A privacy choice is a mechanism by which a user is allowed to control one or more practices associated with the collection or processing of data about them. This may include opting or out of some data practices, requesting to review and/or correct data about oneself, or requesting the deletion of such data and more.

User-Oriented Privacy Threat
A user-oriented privacy threat is any failure to notify the user about the data practices and/or give the user control over their data, or any attempt to manipulate the user into revealing or sharing their personal information.

ii

Introducing AccuFrame

AccuFrame is a fictional online glasses retailer, and recently introduced a "virtual try-on" feature to its mobile application ("VTO tool"), where users are able to see themselves wearing products such as glasses virtually so that they can make more informed purchasing decisions. The VTO tool requires users to allow the app to access their device's camera to scan their face, collect and then analyze their biometric data in order to make personalized product recommendations.

Imagine you are a privacy consultant and are tasked with reviewing a selection of privacy notice and choice interfaces implemented by AccuFrame in terms of their effectiveness and if they reflect users' needs and expectations.

AccuFrame

iii

Description of the user journey

The user persona, Chloe, uses the AccuFrame app to find suitable eyeglasses by using the virtual try-on feature (VTO).

In order to use VTO, Chloe is required to go through a privacy consent page regarding the collection of biometric data and undergo a face scan.

Upon providing consent and completing a VTO quiz to get more personalized frame recommendations, Chloe now wishes to withdraw her consent and exercise control over her biometric data by reviewing the app's privacy policies and settings.

iv

How to read the storyboard

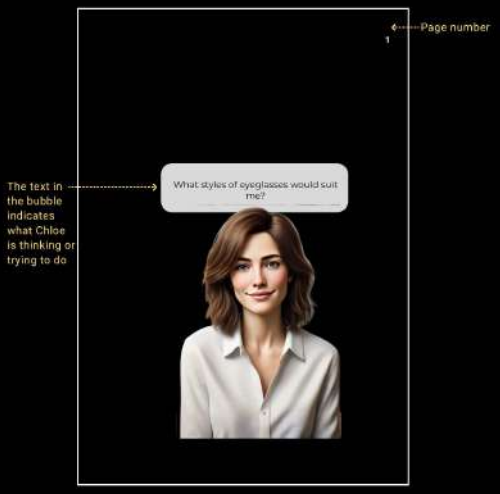

How to read the storyboard

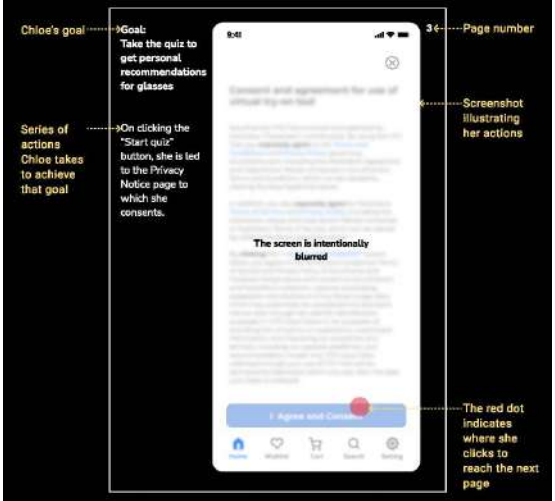

*Note: Chloe's journey to protect her biometric data is illustrated through a series of **static screenshots.***

AccuFrame User Journey

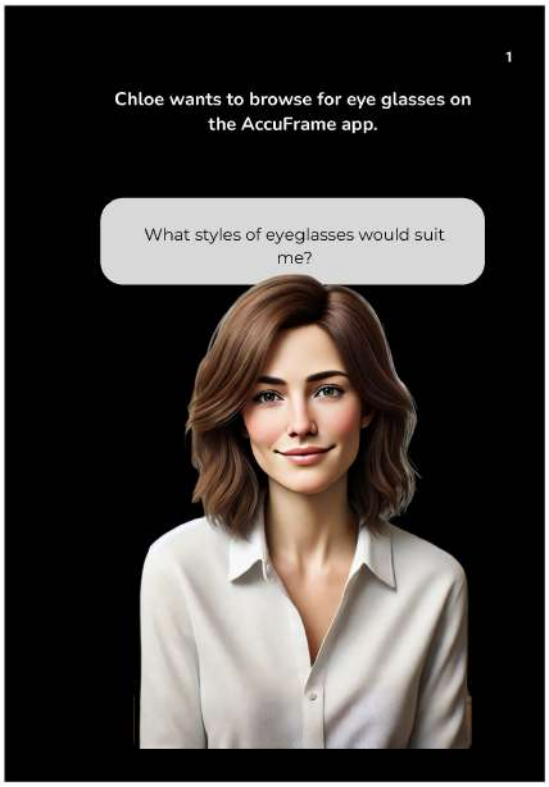

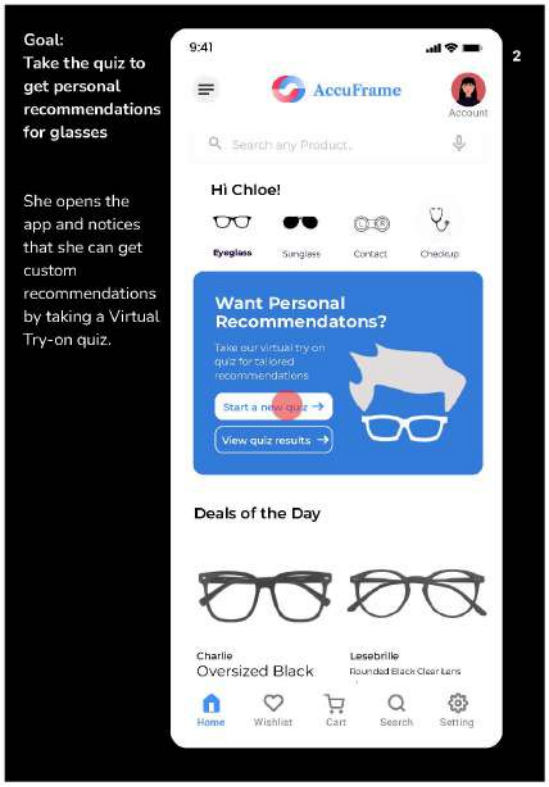

Goal:
Take the quiz to get personal recommendations for glasses

On clicking the "Start quiz" button, she is led to the Privacy Notice page to which she consents.

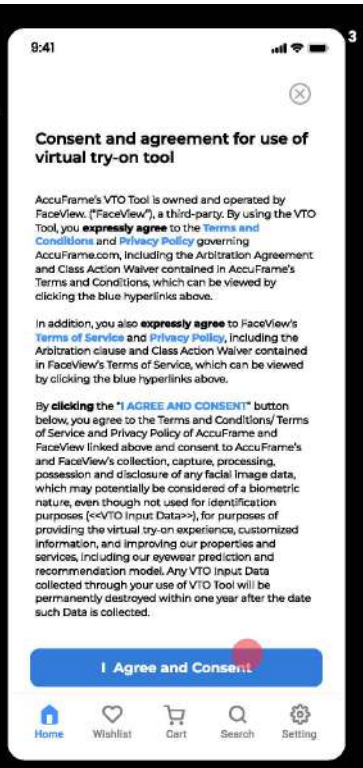

Goal:
Take the quiz to get personal recommendations for glasses

On clicking the "I Agree and Consent" button, she is taken to the face scanner in the Virtual Try-on feature. Then clicks "Next" for Step 2.

Goal:
Take the quiz to get personal recommendations for glasses

After scanning her face, she types in her desired price range and clicks "Next" for Step 3

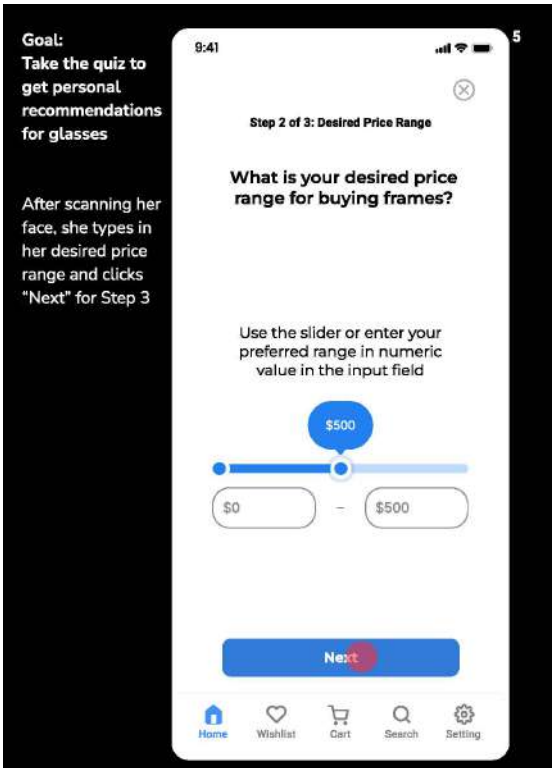

Goal:
Take the quiz to get personal recommendations for glasses

She then indicates her preferences by liking or disliking a series of frames. Then she clicks "Finish" to see her recommendations.

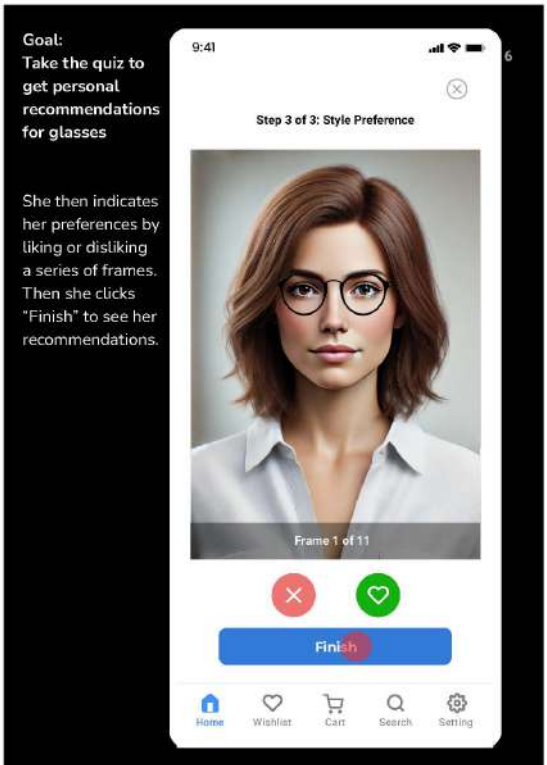

Goal:
Take the quiz to get personal recommendations for glasses

After completing the quiz, she receives frame recommendations based on her face scan, desired price range and preferred styles.

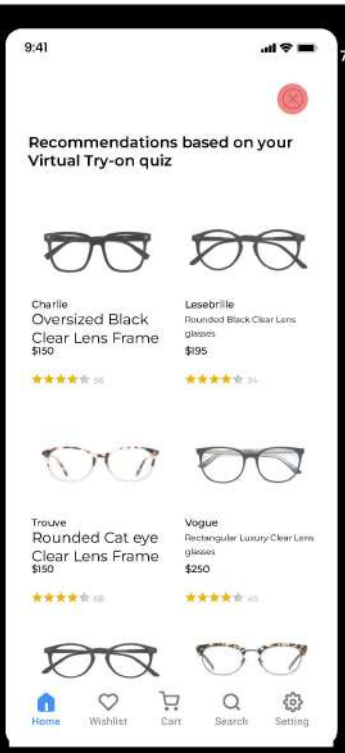

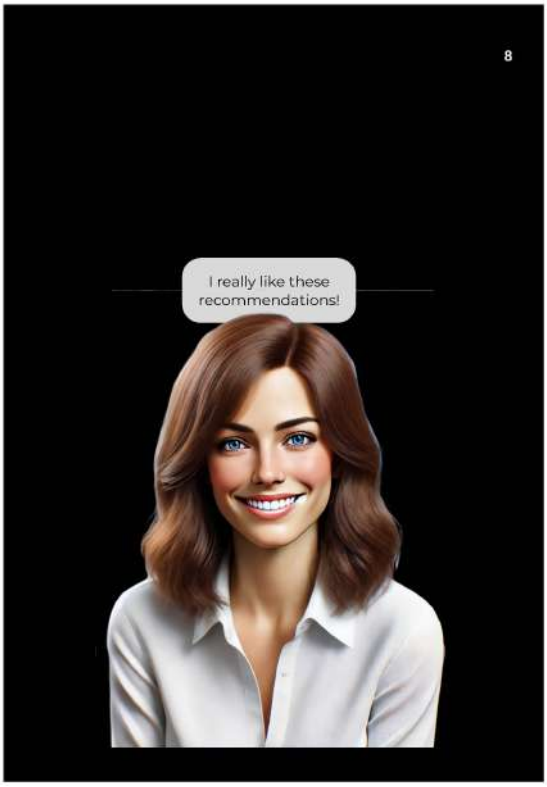

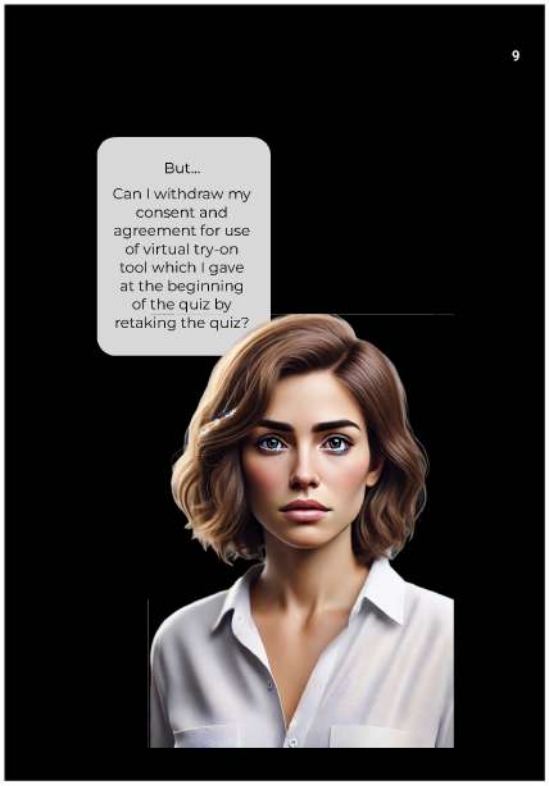

Goal:
Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

After exiting the Recommendations page, she returns to the Home Page and tries to retake the Quiz for customized recommendations to reach the consent page.

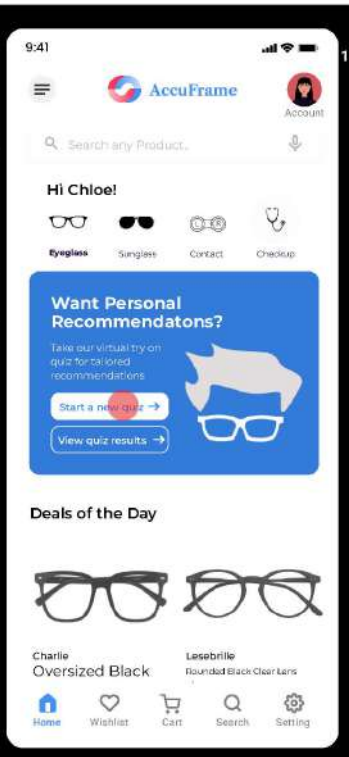

Goal:
Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

On Clicking the "Start a new quiz" button again, she does not see a Consent page this time and is directly taken to the Face Scan page.

11

12

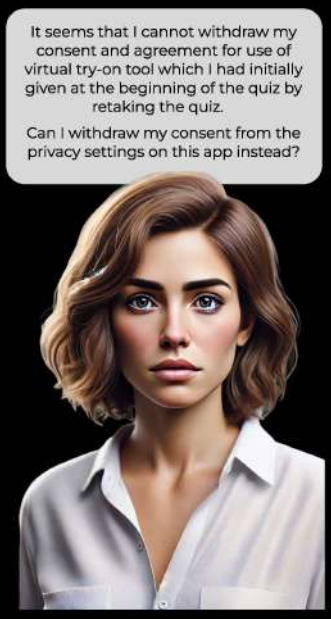

Goal:
Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

Chloe exits the quiz and returns to the Home Screen. She then navigates to the Account menu at the top right of the screen.

13

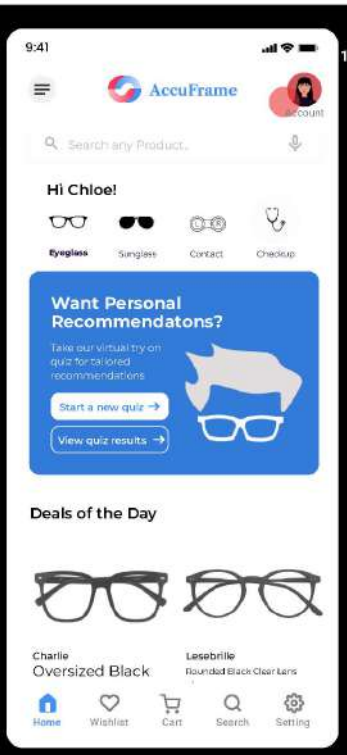

Goal:
Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

She navigates to the Terms and Privacy tab at the bottom of the screen

14

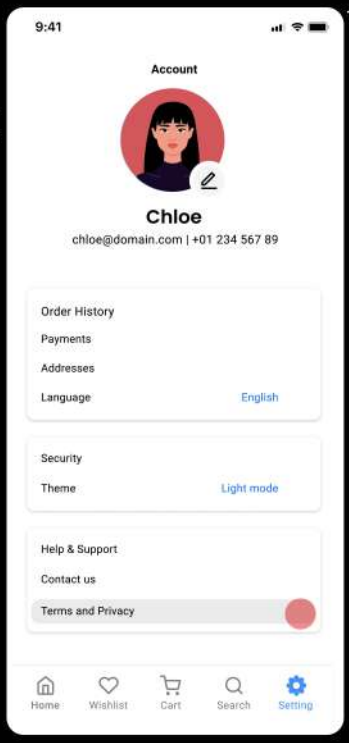

Goal: Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

She navigates to "Your Privacy Choices"

Goal: Withdrawing consent and agreement for use of virtual try-on tool to protect biometric data

She lands on this page, and is not able to find choices about biometric data.

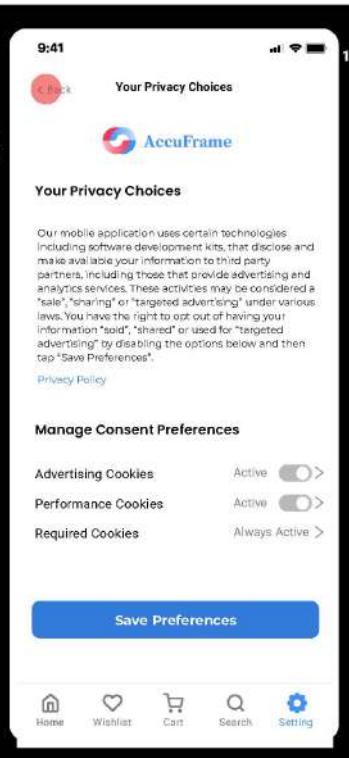

Goal: Deleting biometric data

She can not withdraw her consent regarding the use of her biometric data. So she clicks on the "Privacy Policy and Your Rights" to learn more about how she can delete her data.

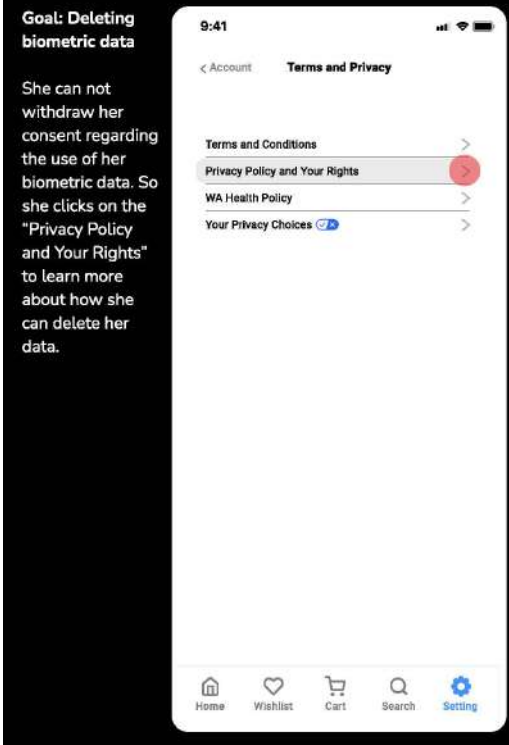

Goal: Deleting biometric data

She scrolls down to read the Privacy policy

You can read the privacy policy in the next page.

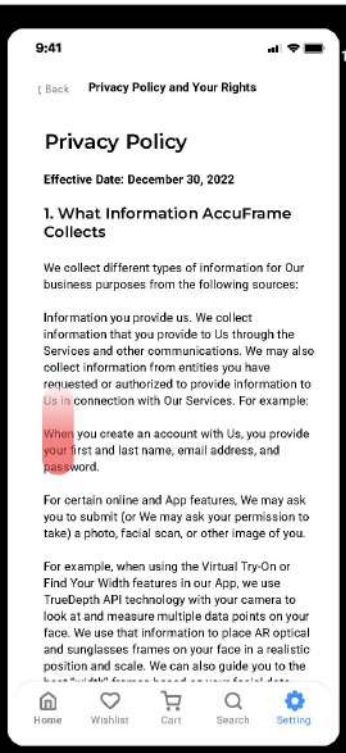

Thank you

You have reached the end of the storyboard

Goal: Deleting biometric data

After she clicks on the Data Request link, she is led to the Data Request page.
She enters the name of the State from which she belongs - Pennsylvania, to see that she has no data rights on the platform.

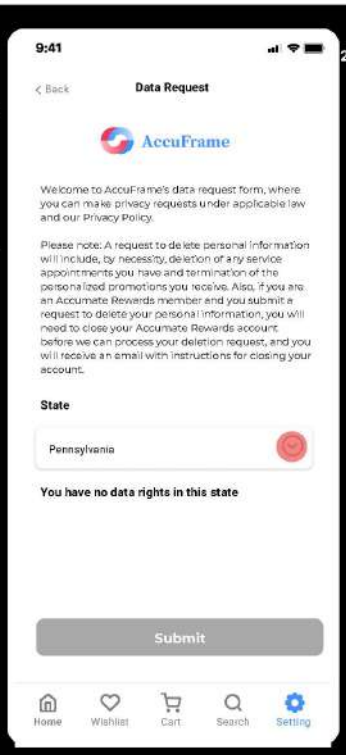

Privacy Policy

Effective Date: December 30, 2022

1. What Information AccuFrame Collects

We collect different types of information for Our business purposes from the following sources:

Information you provide us. We collect information that you provide to Us through the Services and other communications. We may also collect information from entities you have requested or authorized to provide information to Us in connection with Our Services. For example:

When you create an account with Us, you provide your first and last name, email address, and password.

For certain online and App features, We may ask you to submit (or We may ask your permission to take) a photo, facial scan, or other image of you.

For example, when using the Virtual Try-On or Find Your Width features in our App, we use TrueDepth API technology with your camera to look at and measure multiple data points on your face. We use that information to place AR optical and sunglasses frames on your face in a realistic position and scale. We can also guide you to the best “width” frames based on your facial data points. We do not store any of these scans or measurements and we only collect and use that data while you are using the Virtual Try-On feature. We do not share these scans or measurements with any third parties.

2. How We Use Your Information

We may use information We collect to operate Our business, including:

- To provide, personalize, and improve Our products and Our Services
- To conduct internal research
- To deliver content and marketing communications that We think may interest you, including ads or offers tailored to you based upon your browsing and usage history, both within these Services and on other websites and mobile applications
- To map your facial features in order to provide products and recommendations, including to facilitate a Virtual Try-On with augmented reality
- To comply with the law, regulations, and other legal obligations
- To properly verify your identity, prevent fraud and enhance security
- To audit and provide reporting relating to particular transactions and interactions, including online interactions you may have with Us or others on Our behalf

- To allow you to utilize features and personalized content within Our Services when you grant Us access to information from your device, including location-based services
- For short-term, transient use including contextual customization of ads
- For other purposes, as permitted by law or to which you consent

3. How Long We Keep Your Information

We will retain your Personal Information only for as long as we reasonably consider it necessary for achieving the purposes set out in this Privacy Policy, or for as long as we are legally required to retain it.

4. How We Share Your Information

If We share your information, We do so to support Our business, including with:

Affiliates and Service Providers. We've figured out ways to do a lot of things on Our own, but We haven't quite figured out how to do all of it. We may share information with Our corporate entities and affiliates, service providers, data processors, third party contractors, payment processors, others who perform services for Us, such as:

- order fulfillment
- delivery services
- payment processing
- vision insurance claim processing
- account registration
- website-related services, such as web hosting
- improvement of website-related services and features
- maintenance services
- marketing services
- data analytics

Other Entities. We may share information as part of a merger, acquisition, or other sale or transfer of all or part of Our assets or business or with other entities as you have authorized or requested.

With Marketing, Analytics, and Advertising Partners. Some of Our third-party advertising partners use cookies and other technologies to collect information about your online activities on Our Services and across other online services in order to deliver more relevant advertising when you are using the Services or other websites.

To Protect Us and Others. We reserve the right to access, read, preserve, and disclose any information that We reasonably believe is necessary to comply with any applicable law, court order, subpoena, legal process, or enforceable governmental request; cooperate with law enforcement; enforce or apply this Privacy Policy, Our Terms of Use, and other agreements;

detect, prevent, or otherwise address fraud, security or technical issues; or protect the rights, property, or safety of AccuFrame, Our employees, Our users, or others.
The third parties that receive your information are required to treat your Personal Information in accordance.

5. Your Choices

US State Rights

Depending on where you live, you may have the following rights:

Right to Know / Access / Portability. Request that we confirm whether we are processing your information, obtain details about the processing activities, and obtain a copy of such information, subject to exceptions.

Right to Delete. Request that AccuFrame delete your information, subject to exceptions.

Right to Correct. Request that AccuFrame correct your information, subject to exceptions. Regardless of where you live, you may view and correct your account information by logging into your account online or contacting us and requesting such changes.

Limit the Processing of Sensitive Personal Information. If you live in California, you have a right to ask that AccuFrame limit how it processes your sensitive personal information, subject to exceptions.

Right to Withdraw your consent. You may withdraw your consent at any time, subject to legal or contractual restrictions and reasonable notice. However, in some circumstances, we may have to limit the products and services provided to you.

Withdrawing your consent for integral purposes. You may withdraw your consent for purposes that are integral to the provision of our products and services, but then you might not be able to proceed with your intended interactions or transactions with AccuFrame or otherwise receive the full benefit of AccuFrame's products and services.

Withdrawing your consent for additional purposes. You may also withdraw your consent for purposes that are not integral to the provision of our products and services.

Withdrawing your consent for additional purposes that are not integral to the provision of our products and services will not impact the provision of our products and services to you.

Those additional purposes include:

- to communicate with you for the purposes of providing you with advertising and marketing messages pertaining to additional products or services that may be of interest to you;
- to administer and facilitate your participation in promotions related to AccuFrame;
- to conduct surveys on the quality of our products and services.

To withdraw your consent, you could modify your privacy settings or contact us by email.

Opt-Out of Selling, Sharing, or Targeted Advertising. Our use of certain online tracking technologies may be considered a “sale”, “sharing”, or “targeted advertising” under applicable law. You can opt-out of this type of activity by clicking the “Do Not Sell or Share My Personal Information” link at the bottom of our Sites. Because we may also engage in a “sale”, “sharing” or use of your personal information for targeting advertising purposes outside of the context of online tracking technologies as well, you may separately unsubscribe from this process by submitting the request via the link below. Finally, you may exercise your right to opt-out of the online tracking technologies process by using the Global Privacy Control (GPC) (on the browsers and extensions that support such a signal).

If you choose to use the GPC signal, you will need to turn it on for each supported browser extension you use.

Right to Appeal. Appeal any denial of a Right to Access, Delete, Correct, or Unsubscribe, as described above.

To exercise all of your above Choices, please click on the Data Request.

Additional Health Data Rights. If you are a resident of Nevada, Connecticut, or Washington, you have additional rights concerning your health data, including the right to submit a request to know or access, deletion, and the right to appeal any denial of these rights. You may exercise these rights by clicking on the Health Data Request.

Managing Your Preferences & Account

Opt-Out of Promotional Messaging. Regardless of where you live, if you would like to stop receiving promotions, special offers or member-exclusive events, you can update your email preferences by visiting My Account on accuframe.com or on our mobile application. You may also notify our Guest Services team by visiting our Contact Us page. Please note it may take up to 6-8 weeks to stop receiving these communications after updating your preferences.

Opt-Out of Texting. Regardless of where you live, you may unsubscribe from AccuFrame text messages, reply "stop" to text messages sent from 95637 (ACCU). This will unsubscribe you from all AccuFrame text message campaigns from 95637 (ACCU). To unsubscribe from text messages from AccuFrame Messenger service, text STOP to (630) 410-9968. This will opt you out of all AccuFrame Messenger text message campaigns from (630) 410-9968. To opt out of transactional text messages from AccuFrame , text STOP to 46373. This will opt you out of all AccuFrame transactional text messages from 46373. Message and data rates may apply.

Opt-Out of Push Notifications. Regardless of where you live, you may unsubscribe from AccuFrame sending you push notifications by adjusting the permissions on your device.

Opt-Out of Precise Geolocation. Most browsers and mobile devices allow users to enable or disable precise geolocation using pop-ups or controls located in the settings menu. Regardless of where you live, if you have permitted AccuFrame to access your precise geolocation data, you may at any time opt-out from further allowing AccuFrame to access your precise geolocation data by adjusting the permissions in your browser or device.

Delete Your AccuFrame Account on the iPhone. Regardless of where you live, you may delete your AccuFrame account at any time from your AccuFrame application. We will email you instructions to confirm your identity and finalize your deletion after you submit your account deletion request. Please note: deleting your account will result in the loss of any points associated with your loyalty account. AccuFrame may still retain certain information connected to your purchase history as required under applicable law.

To exercise all of your above Choices, please click on the Data Request.

APPENDIX B
STORYBOARD – BEYOND

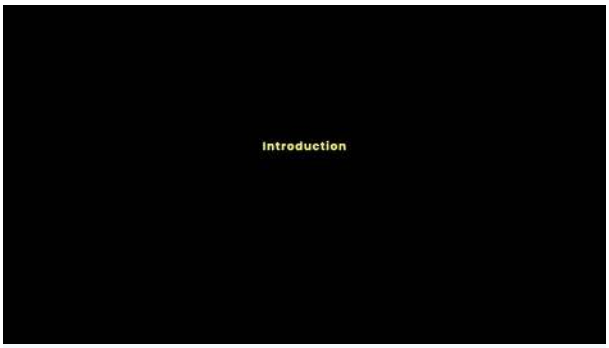

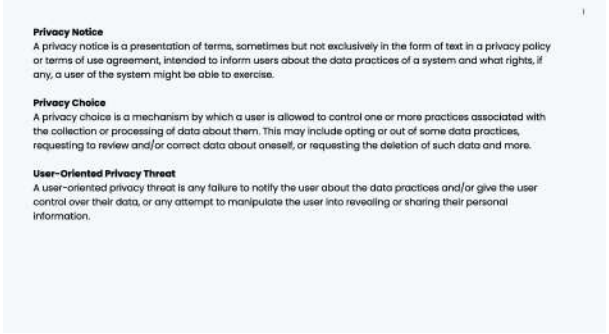

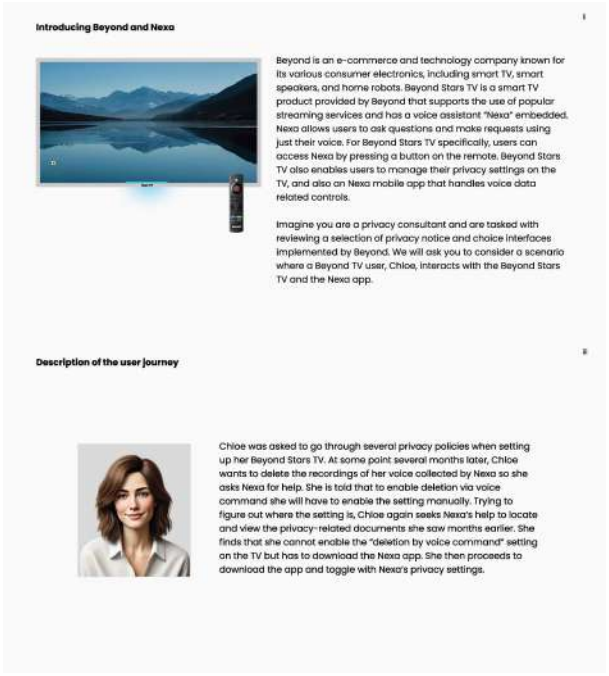

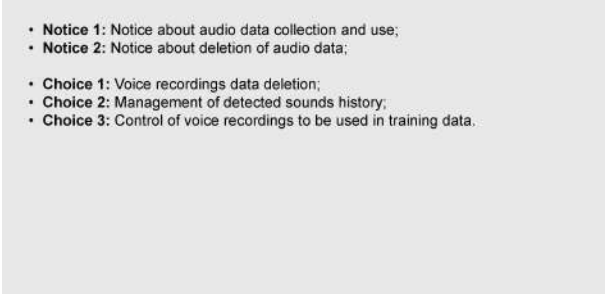

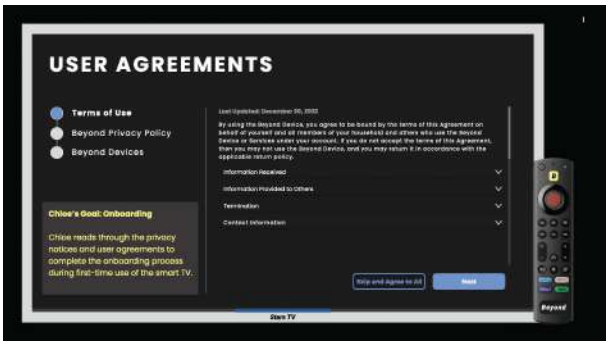
USER AGREEMENTS
Terms of Use
Beyond Privacy Policy
Beyond Devices
Chloe's Goal: Onboarding
Chloe reads through the privacy notices and user agreements to complete the onboarding process during first-time use of the smart TV.

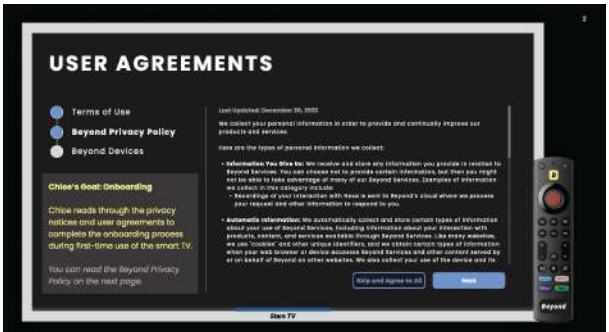
USER AGREEMENTS
Terms of Use
Beyond Privacy Policy
Beyond Devices
Chloe's Goal: Onboarding
Chloe reads through the privacy notices and user agreements to complete the onboarding process during first-time use of the smart TV.
You can read the Beyond Privacy Policy on the next page.

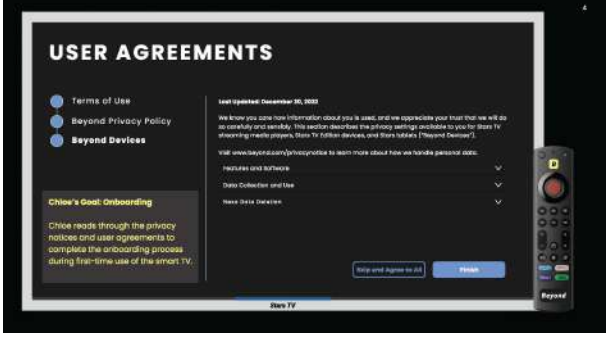
USER AGREEMENTS
Terms of Use
Beyond Privacy Policy
Beyond Devices
Chloe's Goal: Onboarding
Chloe reads through the privacy notices and user agreements to complete the onboarding process during first-time use of the smart TV.


Chloe's Goal: Understanding what voice recordings have been stored from previous interactions
Chloe interacts with the TV through the voice assistant for several months and wonders what voice recordings have been stored from her previous interactions with Nexa.
Hey Nexa, what voice recordings of mine have been stored?
...the following are the latest voice recordings that have been recorded and stored:
"Pause" 5:00 pm July 23, 2024
"Play" 3:30 pm July 23, 2024
"Go to the next channel" 3:29 pm July 23, 2024 ...
You can view all on the Nexa App
Chloe's Goal: Deleting her voice data
Chloe feels that it is unnecessary to store such voice recordings, so she interacts with Nexa to delete those voice recordings.
Hey Nexa, please delete all my recordings.
...the feature of deleting voice data collected via voice commands is currently disabled. Please manually modify your settings to enable that feature...

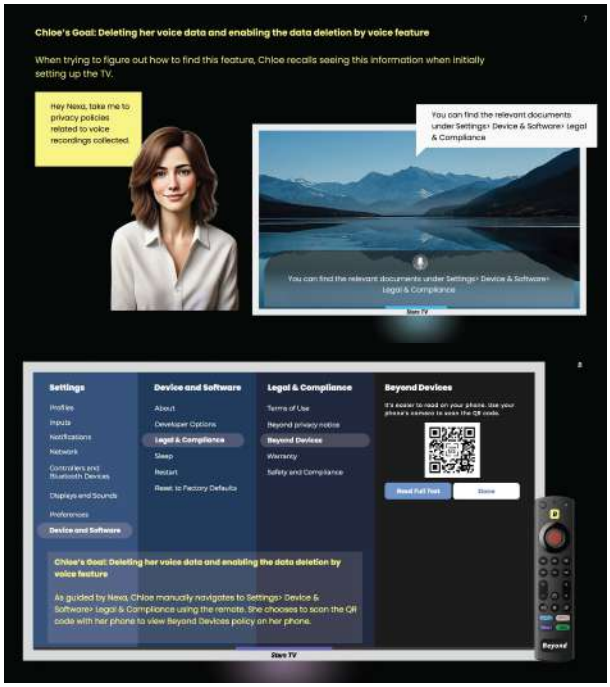
Chloe's Goal: Deleting her voice data and enabling the data deletion by voice feature
When trying to figure out how to find this feature, Chloe recalls seeing this information when initially setting up the TV.
Hey Nexa, take me to privacy policies related to voice recordings collected.
You can find the relevant documents under Settings> Device & Software> Legal & Compliance
Settings
Device and Software
Legal & Compliance
Beyond Devices
Chloe's Goal: Deleting her voice data and enabling the data deletion by voice feature
As guided by Nexa, Chloe manually navigates to Settings> Device & Software> Legal & Compliance using the remote. She chooses to scan the QR code with her phone to view Beyond Devices policy on her phone.


Chloe's Goal: Deleting her voice data and enabling the data deletion by voice feature
She reads the Beyond Device's Privacy Policy.
You can read the Beyond Devices document on the next page.
Beyond Devices
Chloe's Goal: Deleting her voice data
Chloe reads under the "Data Deletion" section that she can delete all her voice recordings by saying, "Nexa, delete everything I said."
Nexa, delete everything I said.
...the feature of deleting voice data collected via voice commands is currently disabled. Please manually modify your settings to enable that feature...

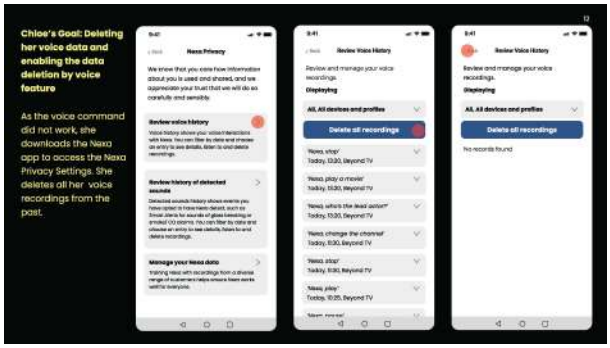
Chloe's Goal: Deleting her voice data and enabling the data deletion by voice feature
As the voice command did not work, she downloads the Nexa app to access the Nexa Privacy Settings. She deletes all her voice recordings from the past.
Nexa Privacy
Review Voice History
Delete all recordings

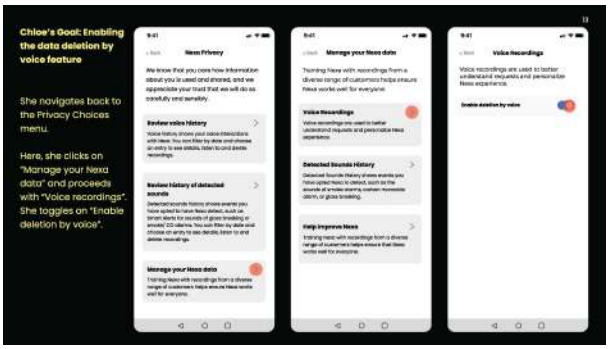
Chloe's Goal: Enabling the data deletion by voice feature
She navigates back to the Privacy Choices menu.
Here, she clicks on "Manage your Nexa data" and proceeds with "Voice recordings". She toggles on "Enable deletion by voice".


Chloe's Goal: Exploring other privacy features
Even though her goal is served, she decides to explore the other settings out of curiosity.
Chloe's Goal: Explering other privacy features
Next, she explores the "Help Improve Nexa" setting. She toggles off "Use of voice recordings" for improving Nexa.


Thank you
This is the end of the user journey

Beyond Privacy Policy

Last Updated: December 30, 2022

We collect your personal information in order to provide and continually improve our products and services.

Here are the types of personal information we collect:

- **Information You Give Us:** We receive and store any information you provide in relation to Beyond Services. You can choose not to provide certain information, but then you might not be able to take advantage of many of our Beyond Services. Examples of information we collect in this category include:
  - Recordings of your interaction with Nexa is sent to Beyond's cloud where we process your request and other information to respond to you.
- **Automatic Information:** We automatically collect and store certain types of information about your use of Beyond Services, including information about your interaction with products, content, and services available through Beyond Services. Like many websites, we use "cookies" and other unique identifiers, and we obtain certain types of information when your web browser or device accesses Beyond Services and other content served by or on behalf of Beyond on other websites. We also collect your use of the device and its features, such as your navigation of the home screen, selection of device settings (such as device language, display size, WiFi and Bluetooth options), or open and close of apps and services.

We use your personal information to operate, provide, develop, and improve the products and services that we offer our customers. These purposes include:

- **Purchase and delivery of products and services.** We use your personal information to take and handle orders, deliver products and services, process payments, and communicate with you about orders, products and services, and promotional offers.
- **Provide, troubleshoot, and improve Beyond Services.** We use your personal information to provide functionality, analyze performance, fix errors, and improve the usability and effectiveness of Beyond Services.
- **Recommendations and personalization.** We use your personal information to recommend features, products, and services that might be of interest to you, identify your preferences, and personalize your experience with Beyond Services.
- **Provide voice, image, and camera services.** When you use our voice, image and camera services, we use your voice input, images, videos, and other personal information to respond to your requests, provide the requested service to you, and improve our services. For more information about Nexa voice services, click here.
- **Comply with legal obligations.** In certain cases, we collect and use your personal information to comply with laws. For instance, we collect from sellers information regarding place of establishment and bank account information for identity verification and other purposes.
- **Advertising.** We use your personal information to display interest-based ads for features, products, and services that might be of interest to you. We do not use information that personally identifies you to display interest-based ads.

In addition, to the extent required by applicable law, you may have the right to request access to or delete your personal information. If you wish to do any of these things, you may go to Data Privacy Queries. Depending on your data choices, certain services may be limited or unavailable.

Beyond Devices

We know you care how information about you is used, and we appreciate your trust that we will do so carefully and sensibly. This section describes the privacy settings available to you for Stars TV streaming media players, Stars TV Edition devices, Stars tablets and Kindle e-readers (“Beyond Devices”).

Visit here to learn more about how we handle personal data.

Features and Software

Your Beyond Device may have features that allow you to access Nexa voice services or otherwise use your voice to perform certain tasks, such as check the weather, add a calendar entry, perform a search, or operate other connected products. When you use voice services, we may process your voice input and other information (such as location) in the cloud to respond to your requests and to improve your experience and our products and services. Your use of Nexa is subject to the Nexa Terms of Use (here). Learn more about Nexa voice services and how it works at here.

Data Collection and Use The operating system of your Beyond Device collects data about your use of the device and its features, such as your navigation of the home screen, selection of device settings (such as device language, display size, WiFi and Bluetooth options), or open and close of apps and services.

We use this data to provide you with the features of your device, for example, to fulfill the actions you take on the device (such as opening an app), to allow you to restore the settings on your device, and to help you troubleshoot when you experience technical issues. We also use it to develop and improve products and features for all our customers and to gain insights into how our products are being used, assess customer engagement, identify potential quality issues, analyze our business, and customize marketing offers.

Nexa Data Deletion We provide ways to manage your data collected regarding your conversation with Nexa while using Stars TV streaming media players, Stars TV Edition devices, Stars tablets that support Nexa to improve your experience. You can delete the voice recording of your last request by saying "Nexa, delete what I just said," delete the voice recordings associated with your account for the day by saying "Nexa, delete everything I said today," or all the voice recordings associated with your account by saying, "Nexa, delete everything I said." In addition, with a voice ID, you can say “Nexa, update my privacy settings” or “Nexa, auto-delete my voice recordings” to have your voice recordings older than 12 months deleted automatically, or to have them not saved at all. To review and delete the voice recordings associated with your account (including any audio resulting from a false wake), you can go to Settings > Nexa Privacy in your Nexa app.

Deleting voice recordings may degrade your Nexa experience. If you choose not to save any voice recordings, voice ID may not work.

If you have changed your default marketplace while using a Nexa-enabled product, you will need to delete all Nexa voice recordings associated with your account separately for each marketplace. To learn how to transfer your Beyond account to another marketplace, go here.

We may still retain other records of your Nexa interactions, including records of actions Nexa took in response to your request. This allows us, for instance, to continue to provide your reminders, timers, and alarms, process your orders, remember the things you've taught Nexa, and show your shopping and to-do lists and messages sent through Nexa Communications. If your request was processed by a Nexa skill, deleting your voice recordings does not delete any information retained by the developer of that skill (skill developers do not receive voice recordings).

APPENDIX C
RECRUITING EMAIL

Hello,

We're a team of researchers from *ANONYMOUS Institution* and we're recruiting privacy professionals to participate in our research study. We're studying ways to help analysts identify user-oriented threats related to privacy notice and choice. A user-oriented privacy threat is a failure to notify the user about data practices and/or give the user control over their data, or an attempt to manipulate the user into revealing or sharing their personal information.

Join us for a 90-minute, hands-on interview session over Zoom between now and *end date*. We'll provide you with a fictionalized scenario and associated notice and choice interfaces and ask you to walk us through how you would identify the user-oriented privacy threats. The results will inform the development of tools to help make this task easier.

You will need to be based in the US in order to participate and should expect to read relatively long texts in English as part of the interview session. You will be given the option to choose either a T-shirt or tote bag to thank you for your valuable time and insights upon the completion of the study. We will mail you your thank-you gift after the study is over.

If you are interested in participating, please fill out our screening survey, which will ask you to provide some information regarding your background. We will get back to you with further information to confirm your participation and schedule a convenient time for the interview. If you have any questions about our research, please feel free to email *ANONYMOUS*!

Link to screening survey: *ANONYMOUS link*

Best Regards,

ANONYMOUS

APPENDIX D
SCREENING SURVEY

We are a research team from *ANONYMOUS Institution* looking to recruit participants to join an interview study to enhance privacy notice and choice in user interactions. If you: 1) are at least 18 years old; 2) are currently based in the US; 3) are able to read relatively long texts in English; and 4) have either academic or industrial experience in the field of privacy, we would like to hear from you! Please answer each question carefully, as your responses are crucial for determining your eligibility for our study. If selected, you will be invited to a study that will take around 90 minutes. We truly appreciate your time and efforts. Thank you!

Q1. What is your preferred name? *(full response filled)*

Q2. What is the best email address we should use to contact you throughout your participation including scheduling and payment purposes? *(full response filled)*

Q3. Which of the following do you identify yourself with?

- Female
- Male
- Non-binary
- Prefer not to answer

Q4. What is your age? *(full response filled)*

Q5. What is your highest level of education?

- High School Diploma
- Associate degree
- Bachelor's degree
- Master's degree
- Professional degree (e.g., MD, JD)
- Doctorate (e.g., PhD, EdD)
- Prefer not to answer

Q6. Can you describe your current occupation? If you are currently a student, can you please share the school and program you are in? *(full response filled)*

Q7. Please tell us about your experience in privacy briefly. Note that we are using this question to confirm whether you are eligible for our study so please be as thorough as you can. *(full response filled)*

APPENDIX E
INTERVIEW SCRIPT

A. Introduction

Hello, and thank you for participating in our study. My name is [], and I'll be your interviewer today. Joining me is [], who will be responsible for taking notes. We are part of a research team that focuses on finding ways to help analysts identify and categorize user-oriented privacy threats in the design of products and services. I want to assure you that all the information we collect today will be kept confidential. At any point in time during the interview, if you want to terminate your participation, please let us know. If you participate until the end of this interview, we will send you a T-shirt (or a tote bag) to the mail address you provide. We will be recording this session to ensure we accurately capture your feedback and thoughts. In this interview study, we will send you two documents, and ask you to open these two documents in your browser and share screens. If you have questions or concerns about screen sharing, including how it works, please let us know now, and we will demonstrate how screen sharing works. Thank you for signing the consent form and filling out the survey questions from earlier. Just to confirm, do we have your consent to record and go ahead with the interview? [Start screen recording] Thank you for giving your consent to record, and we will start the interview now.

B. Background and General Questions

- **Q1** Can you share your experience in the field of privacy? How long have you been working in the field and what are the areas you specialized in?
- **Q2** What's your experience with privacy notice and choice in digital interfaces? Say, have you ever participated in designing/evaluating one in terms of its effectiveness? (If no was answered to the first part of this question, ask about participants' experience interacting with notice and choice in daily lives)
- **Q3** Have you ever heard of privacy threat modeling? Can you give some examples?

- **Q4 (No Taxonomy Group Only)** As mentioned in the recruitment, we encourage you to bring any privacy threat modeling framework as you see fit with you for today's interview. Have you brought anything with you to use today?

C. Identify Threats

1) Introducing the scenario: Participants will be randomly assigned to one of the two scenarios described below.

- *AccuFrame*

We want to introduce a fictitious company, AccuFrame, which is an online glasses retailer, and recently introduced a "virtual try-on" feature to its mobile application. This feature allows users to see themselves wearing products such as glasses virtually when using the app so that they can make more informed purchasing decisions.

You will be presented with a specific user interaction scenario, which includes a series of actions taken by a user of AccuFrame, Chloe, when she uses the virtual try-on feature on AccuFrame and navigates through privacy notices and choices mechanisms implemented by AccuFrame. For this interview, you are asked to act as a privacy analyst, to examine Accuframe, and to identify possible privacy threats focusing on the implemented set of specific notices and choices.

- *Beyond*

We want to introduce a fictitious company, Beyond, which is an e-commerce and technology company known for its various consumer electronics. Beyond has a smart TV product that supports the use of popular streaming services, and the smart TV includes a voice assistant, Nexa, which allows users to ask questions and make requests using their voice. Nexa also has a mobile application that handles voice data-related controls.

You will be presented with a specific user interaction scenario, which includes a series of actions taken by a user of Beyond, Chloe, when she initializes the smart TV setup, interacts with the voice assistant, and navigates through privacy notices and choices mechanisms implemented by Beyond and Nexa. For this interview, you are asked to act as a privacy analyst to examine Beyond and Nexa and identify possible privacy threats focusing on the implemented set of specific notices and choices.

2) Going over the scenario (No Taxonomy Group): I will now send you the link to the scenario PDF. I will also send you a Google Doc so you can record all the threats you identify in the scenario. Please share your screen once you open these two documents. Also, if you bring any framework with you today, please feel free to open it as well.

(Explaining and walking participants through the 2 documents) Please first take a look at the threat list doc. The first section of this document includes some information on your task as a participant as well as some instructions on how to fill out the table when you are recording the privacy threats. Please take some time to read this section and let me know if you have any questions.

Now please switch to the scenario PDF. The first few pages provide some information and instructions that may help you better understand the task and this storyboard. Please take some time to go over the first several pages.

Please take about 5 minutes to quickly skim through the remaining pages of the scenario pdf to familiarize the actions that Chloe takes.

Again, you are asked to put threats you identified regarding the 4 (or 5) notices and choices in different tables. You can use the table of contents on the left to jump directly to the relevant table.

Now, please take some time to go through the scenarios and fill out the tables per the threat list's instructions. As you type things down, please verbally describe your thoughts. For the importance rating, please also provide your reasoning. During this process, please feel free to use any privacy threat modeling framework or any other tools to help you identify the threats in the scenario.

3) Going over the scenario (With Taxonomy Group): I will now send you the link to the scenario PDF. I will also send you a Google Doc so you can record all the threats you identify in the scenario. Please share your screen once you open these two documents. Also, if you bring any framework with you today, please feel free to open it as well.

(Explaining and walking participants through the 2 documents) Please first take a look at the threat list doc. The first section of this document includes some information on your task as a participant as well as some instructions on how to fill out the table when you are recording the privacy threats. Please take some time to read this section and let me know if you have any questions.

So as you have just read, you are asked to put threats you identified regarding the 4 (or 5) notices and choices in different tables. You can use the table of contents on the left to jump directly to the relevant table.

All 4 (or 5) tables include the exact same set of threats. Please go through the threats listed in the first table to familiarize yourself with the threats and their definitions. If there are any threats that you are not entirely sure of their meaning, please click on the link to see more specific details and definitions. Also feel free to let me know if you have any questions. (If participants go on to the other tables, tell them that the threats listed are the same, the other tables are just copies of the first one such that they have one set of tables per notice or choice)

Now please switch to the scenario PDF. The first few pages provide some information and instructions that may help you better understand the task and this storyboard. Please take some time to go over the first several pages.

Please take about 5 minutes to quickly skim through the remaining pages of the scenario pdf to familiarize the actions that Chloe takes.

(Have them switch back to the threat list doc) Again, you are asked to put threats you identified regarding the 4 (or 5) notices and choices in different tables. You can use the table of contents on the left to jump directly to the relevant table.

Now, please take some time to go through the scenarios and fill out the tables per the threat list's instructions. As you type

things down, please verbally describe your thoughts. For the importance rating, please also provide your reasoning.

D. Follow-up Discussion

- **Q1** (No Taxonomy Group) What was your process for analyzing these storyboards to identify threats? [If the participant mentions bringing some privacy threat modeling framework] Did you use any privacy threat modeling framework as a tool?
- **Q2** (No Taxonomy Group) How easy or difficult was it to identify user-oriented privacy threats in this scenario? Why was it [easy/difficult]?
- **Q3** (No Taxonomy Group) Is there anything that you wish you had during the threat identification process that may be helpful to carry out your analysis?
- **Q1** (With Taxonomy Group) How's your experience using the taxonomy? How easy or difficult was it to identify privacy threats using the taxonomy we provided?
- **Q2** (With Taxonomy Group) How did the taxonomy influence your approach to identifying threats, if at all? (If the answer is somewhat positive) Can you provide specific examples where the taxonomy helps you identify a threat that you might not have considered otherwise?
- **Q3** (With Taxonomy Group) Do you have any suggestions for any improvements or changes to the taxonomy we provided?
- **Q4** Is there anything else you want to share for your threat-identification experience?

(Closure) As we come to the end of our session, I'd like to take a moment to thank you for your time today sincerely. *(Explain the goal of the study and participants being assigned to the no taxonomy/with taxonomy group.)* Your contribution is incredibly valuable to our research, and we're grateful for the perspectives you've provided. Before we conclude, do you have any questions about the study, our research, or anything else we discussed today? Based on your selection, we will send the T-shirt or the tote bag (depending on the participant's answer) to the mailing address you provided, and you will receive an email once it has been shipped. We truly appreciate the time and effort you've put into today's session. Have a wonderful day! I will stop recording at this point.

APPENDIX F
INSTRUCTION FOR NO TAXONOMY PARTICIPANTS

Below is the version of instruction provided to AccuFrame no-taxonomy participants. In the document given to participants, we include an empty table for each of the notice and choice for them to use for threat identification.

What is my task as a participant?

You are asked to act as a privacy analyst, to examine a fictional mobile app called *Accuframe,* and to identify possible shortcomings in the way in which a set of specific notices and choices are implemented in the Accuframe app. The app helps users browse eyeglass frames and uses images of the user's face to show them how the frames would look on them. Please use this Google doc to record all the privacy notice and choice threats you identify when you review Chloe's journey when using AccuFrame.

Here is a list of privacy notices and choices that Chloe will come across in the scenario, which we also want you to focus on during your analysis:

- **Choice 1: Consent to the collection and use of biometric data;**
- **Choice 2: Cookie management options;**
- **Notice 1: Notice about data deletion rights;**
- **Choice 3: Data deletion control, namely ability to request the deletion of one's data.**

Below we provide a copy of 4 tables, one for each of the notice or choice for you to fill in when you review the scenario. Please identify as many user-oriented privacy threats as you can:

- In **column A**, record the page number where you identify the threat (you can find the page number on the top right corner of each page).
- In **column B**, put a few words to describe the threat.
- In **column C**, state what evidence you find for the threat (e.g., don't just say that a particular interface has usability problems, briefly explain the type of problem and point to something you observed that illustrates this).
- In **column D**, here you are requested to report on the importance of the privacy threat as you would as a privacy analyst working for Accuframe. Use a scale of 1 to 5, where 1 indicates "not important at all" and 5 indicates "extremely important".
- In **column E,** put any mitigation suggestions or ideas for alternative designs for a specific threat.
- Feel free to add more rows (by clicking "insert one row below") if needed or leave the rows blank if you can not find any more threats.

APPENDIX G
USERSFIRST THREAT TAXONOMY & INSTRUCTIONS FOR WITH-TAXONOMY PARTICIPANTS

Below is the version for AccuFrame with-taxonomy participants. In the document given to participants, an empty row was added beneath each threat type, and for every notice and choice, a corresponding table, along with an additional threat list table was included.

What is my task as a participant?

Some companies would like to try to improve their notice and choice experiences but it's not always easy for them to figure out the way to do it. This document describes a privacy threat modeling framework for notice and choice intended to help an analyst systematically assess and refine notice and choice interfaces with the goal of improving usability while complying with relevant regulations. It also details how you are supposed to use this framework as part of the task assigned to you in this study.

You are asked to act as a privacy analyst, to examine a fictional mobile app called *Accuframe,* and to identify possible privacy threats in the way in which a set of specific notices and choices are implemented in the Accuframe app.

You are requested to consider a set of different contexts in which a user persona, Chloe, might interact with this app and determine to what extent a specific set of notices and choices (detailed below) is adequately supported in the current implementation, as captured in a set of screenshots. The screenshots are intended to capture the different ways in which Chloe might interact with the app. For the purpose of this task, you should limit yourself to only the screenshots and tasks that are detailed in the AccuFrame scenario PDF. In other words, you should not speculate about other possible screens not shown in the scenario PDF.

The specific set of 4 notices and choices you are required to analyze are:

- **Choice 1: Consent to the collection and use of biometric data;**
- **Choice 2: Cookie management options;**
- **Notice 1: Notice about data deletion rights;**
- **Choice 3: Data deletion control, namely ability to request the deletion of one's data.**

As you analyze these 4 notices and choices based on the specific screens shown to you, please remember that elements related to these notices and choices may be present in different parts of the interface (e.g., different statements, different options shown in different screens). As you will see in the description of the taxonomy of threats introduced below, some threats can be analyzed by looking at the particular way in which a given notice or choice is implemented in a given screen. Some other threats, however, require taking a broader perspective such as looking for possible inconsistencies in statements made about a given notice or choice in different screens.

Taxonomy of Threats in the Implementation of Privacy Notices and Choices

The taxonomy of threats is organized around four high level categories:

- **Discovery and Use (DU.x):** Threats when it comes to supporting the discovery and practical use of privacy notices and choices
- **Comprehension (C.x):** Threats related to the comprehension of privacy notices and choices
- **Appropriate Choices (AC.x):** Threats in the organization or presentation of choices to users

- **Manipulative Elements (M.x):** Manipulative interfaces such as manipulative statements or manipulative presentation

Using the Tables to Complete your task

Below you will find 4 copies of a table detailing the taxonomy of threats you are requested to use. **Each copy is to be used for one of the four notice and choices you are requested to analyze and is organized around the four top level categories of threats introduced above.** Below each threat, there's one row for you to enter information regarding this threat as indicated below:

- Please use **column A** to record the page number where you identify the threat (you can find the page number on the top right corner of each page).
    - If you identify the same threat in multiple places, please enter all the page numbers where the threat is present.
- In **column B**, state what evidence you find for the threat (e.g., don't just say that a particular interface has usability problems, briefly explain the type of problem you have identified. Please be specific).
- In **column C**, here you are requested to report on the importance of the threat as you would as a privacy analyst working for Accuframe. Use a scale of 1 to 5, where 1 indicates "not important at all" and 5 indicates "extremely important".
- In **column D,** indicate how you would recommend mitigating the threat you have identified.
- For each of these 4 notices and choices, we have added **a table labeled "Additional Threat List."** Please use that table to record any threat you identify in the PDF that does not seem to fall under any of the categories in the threat taxonomy.
- To see a more detailed description of a specific threat, click on the blue links provided in the table. These more detailed descriptions include evaluation questions of that threat along with some practical examples.

Main Menu

The following menu enables you to jump directly to the tables you are requested to use to record your analysis of each notice/choice.

- **Choice 1: Consent to the collection and use of biometric data;**
- **Choice 2: Cookie management options;**
- **Notice 1: Notice about data deletion rights;**
- **Choice 3: Data deletion control, namely ability to request the deletion of one's data.**

Choice 1:Consent to the collection and use of biometric data

Back to menu

A.Page Number	B.Evidence	C.Importance	D.Design Suggestion
Discovery and Use (DU)			
[DU.1] **Nonexistent or Difficult to Locate** Privacy notices and choice mechanisms that are missing or placed in a way that make it difficult for users to locate or be aware of their presence.			
[DU.2] **Ineffective Timing** Privacy notice or choice mechanisms presented at inopportune times that reduce their effectiveness. This includes situations where the privacy policy is the only notice available.			
[DU.3] **Ineffective Channel** Privacy notices or choices delivered through a channel (e.g. website, mobile app, interface built into device) that is neither the primary channel that users use to interact with the system/service nor a reasonably convenient alternate channel wherever the primary channel is not feasible.			
[DU.4] **Lack of Centralized Management** No centralized location (i.e., a privacy dashboard) where users can access and manage all privacy notices and choice mechanisms.			
[DU.5] **Decoupled Notice and Choice** Privacy notices presented to users without direct, convenient, or assisted access to associated privacy choice mechanisms.			
[DU.6] **Poor Organization** *[DU.6.1]* **Lengthy Text that Lacks Structure or Effective Navigation Aids** Privacy notices or choice mechanisms that contain unnecessarily lengthy descriptions while being poorly structured and lack effective navigation aids. *[DU.6.2]* **Too Much Effort to Access Necessary Information (links or layered policy)** Privacy notices and choice mechanisms that are difficult or time-consuming to navigate due to the need for users to follow multiple links, navigate through multiple layers, or expand multiple nodes. *Please specify which of these two threats you are referring to.*			
[DU.7] **Poorly Formatted Notices and Choices** Privacy notices and choice interfaces that cause unnecessary difficulty for users seeking general or specific information due to poor formatting.			
[DU.8] **Dysfunctional components (links, buttons, switches, etc)** Privacy notices or choice mechanisms that include components that do not function as indicated.			
[DU.9] **Distracting Visual/Audio Effects** Privacy notices and choice interfaces that include distracting visual or audio features that could reasonably distract users.			

A.Page Number	B.Evidence	C.Importance	D.Design Suggestion
Comprehension (C)			
[C.1] **Contradictory Statement(s) or Implementation(s)** *[C.1.1]* **Conflicting Statement(s)** Conflicting statements within the privacy notice and choice mechanisms about the same data practice; this includes statements that users are likely to interpret as conflicting even if they can also be interpreted in a way that is not conflicting. *[C.1.2]* **Mismatched Notice Statement and Choice Implementation** Statements in privacy notices regarding available choices and how they are implemented are inconsistent with the service/device's actual choice implementation. *Please specify which of these two threats you are referring to.*			
[C.2] **Inconsistent Terminology** Different terms used throughout the notice and choice interface to describe the same concept or data type.			
[C.3] **Difficult to Understand** *[C.3.1]* **Unclear Terms, Statements, or Choices** Privacy notices and choice mechanisms containing unclear words, phrases, or statements that could lead to confusion, ambiguity, unclearness, or multiple interpretations. *[C.3.2]* **Use of Legal or Technical Jargon** Privacy notices and choice interfaces that use jargon or acronyms that could make it challenging for the intended audience to understand the content without providing informative and non-disrupting explanations. This includes both technical terms that require an expert level of knowledge and legal terms that may not be familiar to most people. *[C.3.3]* **Use of Complex Sentences** Privacy notices and choice interfaces that use language that may be challenging for the intended audience to understand due to the use of long or complex sentence structures or uncommon words. *Please specify which of these three threats you are referring to.*			
[C.4] **Consequences not adequately explained** *This threat only applies to choice* The consequences of privacy choice options are not clearly explained to users in their presented context.			
[C.5] **Inadequate Feedback** *This threat only applies to choice* Privacy choice mechanisms provide none or insufficient feedback in terms of whether the privacy settings have been successfully updated after users submit their choices or info regarding the current state of privacy settings.			
[C.6] **Confusing Buttons/Toggles/Checkbox** *This threat only applies to choice* Choice mechanisms that are presented in a confusing way resulting in user uncertainty as to which state represents each choice.			
Appropriate Choices (AC)			
[AC.1] **Limited Choice** *This threat only applies to choice*			

A.Page Number	B.Evidence	C.Importance	D.Design Suggestion
Privacy choice mechanisms that lack or fail to adequately cover privacy choice options that are required by applicable law or expected by users.			
[AC.2] **Excessive or Redundant Choice Options** *This threat only applies to choice* Privacy choice mechanisms provide too many choices or require too much effort for users to make effective decisions or exercise certain privacy rights.			
[AC.3] **Inadequate or Excessive Granularity** *This threat only applies to choice* Privacy choice options that either fail to encompass user expectation of choice (inadequate granularity) or present too many fine-grained choices (excessive granularity), rendering it unsatisfying or confusing for users when making choices			
[AC.4] **Difficult to Modify Previous Choices** *This threat only applies to choice* Privacy choice mechanisms that make it difficult or impossible for users to modify their choices after submitting the choice to the system.			
Manipulative Elements (M)			
[M.1] **Manipulative Statements** Privacy notices and choice interfaces that use subtle language to manipulate users into taking less privacy-protective actions.			
[M.2] **Visually Manipulative Design** *This threat only applies to choice* A deceptive/dark pattern where the interface encourages users to take invasive privacy actions by using particularly enticing or noticeable font or button colors, different font or button sizes, or manipulative bundling and layouts.			
[M.3] **Asymmetric Effort required for Different Privacy Protection Levels** *This threat only applies to choice* A deceptive/dark pattern in which users need to take more steps for more privacy-protective actions than for less privacy-protective actions.			
[M.4] **Less Privacy Protective Defaults** *This threat only applies to choice* Privacy settings default to options with lower level of protections on privacy.			
[M.5] **Unexpected Choice Alteration** *This threat only applies to choice* User choices that lead to unexpected consequences, especially with regard to other choices.			

Back to menu

Choice 1: Additional Threat List - Use to record threats not listed in the above table

(A)Threat	(B)Page Number	(C)Evidence	(D)Threat Importance	(E)Design Suggestions
			-	

(A)Threat	(B)Page Number	(C)Evidence	(D)Threat Importance	(E)Design Suggestions
			-	
			-	
			-	

Back to menu

More Details on threats from above Tables

Discovery and Use (DU.x)

Threats in this category are related to the discovery and efficient use of privacy notices and choice interfaces.

[DU.1] Nonexistent or Difficult to Locate
Definition:

- Privacy notices and choice mechanisms that are missing or placed in a way that make it difficult for users to locate or be aware of their presence.

Evaluation Questions:

- After typical use patterns, do average users **remain unaware of the presence** of certain privacy notices or choice mechanisms?
- Do the users **find it challenging to deliberately locate** certain privacy notices or choice mechanisms?

Examples:

- **Effective**
  - A website homepage that places links to privacy notices and choice mechanisms in regularly trafficked locations in the form of clearly legible text and effective icons.
  - A mobile application with multiple data collection practices attempts to reduce user effort in seeking privacy notices and/or choice mechanisms regarding a particular data practice through a well organized FAQ section, or other design.
- **Ineffective**
  - An IoT device based data collection system provides privacy notices to users (including incidental users) through inconspicuous and/or poorly labeled QR code.
  - A mobile app putting privacy controls in a settings tab named “General” instead of using more intuitive names such as “Privacy” or “Data Controls.”

[DU.2] Ineffective Timing
Definition:

- Privacy notice or choice mechanisms presented at inopportune times that reduce their effectiveness. This includes situations where the privacy policy is the only notice available.

Evaluation Questions:

- Does the design of timing for privacy notice or choice mechanisms **impede the user's capacity** to pay attention or comprehend important details included in privacy notice or choice?

Examples:

- **Effective**
  - A mobile based location based service offers users the option to enable just-in-time notifications for privacy notice and choice mechanisms, allowing users to make privacy decisions in the actual context of the service's use.
- **Ineffective**
  - All privacy choices for a mobile application, including those that apply only to rarely-used features, are presented to users during the app installation process.

[DU.3] Ineffective Channel

Definition:

- Privacy notices or choices delivered through a channel (e.g. website, mobile app, interface built into device) that is neither the primary channel that users use to interact with the system/service nor a reasonably convenient alternate channel wherever the primary channel is not feasible.

Evaluation Questions:

- Do users interacting with a service or device have to **switch to other channels inconveniently** when they want to change their privacy settings?

Examples

- **Effective**
  - An IoT smart doorbell device with no screen or speakers (or other means of conveying or receiving information) might provide a QR code for authorized users to access privacy notice and choice mechanisms on mobile devices.
  - Privacy notices are provided via public channels if the user's identity is unknown, such as public notices for surveillance cameras.
- **Ineffective**
  - An e-commerce website updates its privacy policy, but only informs registered members through an email notification.
  - A smart speaker that cannot accept voice commands to convey privacy notice information or receive/enact privacy preferences from user voice commands.

[DU.4] Lack of Centralized Management

Definition:

- No centralized location (i.e., a privacy dashboard) where users can access and manage all privacy notices and choice mechanisms.

Evaluation Questions:

- Does the user **need to visit multiple locations** to access information on data practices or submit their privacy preferences for a specific system/service?

Examples:

- **Effective**

    - A mobile application implements a centralized interface that gathers all privacy-related content (either directly or through clearly labeled links) related to its different data collection and use practices, including privacy notices and choice mechanisms.
- **Ineffective**
    - The system implements a privacy notice page, multiple privacy policy pages, and some extra pages detailing state privacy laws in a scattered and disconnected manner.

[DU.5] Decoupled Notice & Choice
Definition:
- Privacy notices presented to users without direct, convenient, or assisted access to associated privacy choice mechanisms.

Evaluation Questions:
- **Notice-Choice Alignment**: Are users presented with a set of corresponding choices with each notice of data practices that users should be able to configure?
- **Choice Accessibility from Notice**: Are choices easy to find from the corresponding sections in the notice? (i.e., are choice interfaces embedded within the notice or are there direct links users can click on in the notice that can successfully lead them to corresponding choices?)

Examples:
- **Effective**
    - Users are informed that they can customize their privacy preferences at will while reviewing the privacy notices, and find the associated privacy choice mechanism to be either directly provided or made available through a clearly labeled link.
- **Ineffective**
    - A privacy notice for a location-based service is delivered to users whenever the service requests location from the user's device; however, the notice does not include direct or convenient access to associated privacy choices such as disabling location tracking or lowering location sharing granularity.

[DU.6] Poor Organization

[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids
Definition:
- Privacy notices or choice mechanisms that contain unnecessarily lengthy descriptions while being poorly structured and lack effective navigation aids.

Evaluation Question:
- Do the privacy interfaces contain **large chunks of texts** that make it difficult for users to extract useful information relating to particular data practices or to particular legal jurisdictions?
- Where appropriate, does the privacy interface layout use **clear headings, bullet points, table of contents, or other visual aids** to facilitate information retrieval?

[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)
Definition:

- Privacy notices and choice mechanisms that are difficult or time-consuming to navigate due to the need for users to follow multiple links, navigate through multiple layers, or expand multiple nodes.

Evaluation Questions:

- Does the privacy notice include **an excessive number of links or expand/collapse buttons** to the degree that users are incapable of extracting useful information without actually visiting the links?
- Does the system/service require users to go through an **unreasonable number of steps** (e.g., expanding all buttons**)** to access specific information?

Examples:

- **Ineffective:**
  - A paragraph in the notice section contains two layers of expand/collapse buttons and five links that lead to different privacy policies (as shown in the figure below).

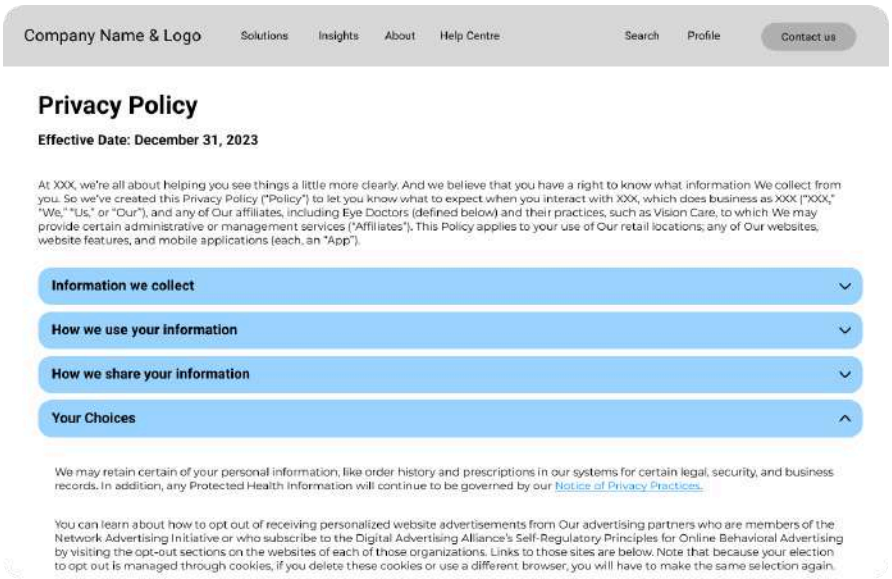

[DU.7] Poorly Formatted Notices or Choice Mechanisms

Definition:

- Privacy notices and choice interfaces that cause unnecessary difficulty for users seeking general or specific information due to poor formatting.

Evaluation Questions:

- Is the **font size, style, and color scheme** easy to read?
- Is the privacy notice/choice optimized for **both desktop and mobile users** in terms of design and formatting?

Examples:

- **Ineffective:**

- Users view the webpage using their mobile phones, only to find that the table on the webpage doesn't format properly, and that they need to scroll both horizontally and vertically to read the full text (as shown in the figure below).

[DU.8] Dysfunctional Components

Definition:

- Privacy notices or choice mechanisms that include components that do not function as indicated.

Evaluation Questions:

- Do all components including links, buttons, and switches **work property and serve their intended purposes**?

Examples:

- **Ineffective**:
  - Clicking into a provided link results in a 404 page not found.

[DU.9] Distracting Visual/Audio Effects

Definition:

- Privacy notices and choice interfaces that include distracting visual or audio features that could reasonably distract users.

Evaluation Questions:

- Are users able to read privacy notices and make privacy choices without being disturbed by any designs implemented by the system?

Examples:

- **Ineffective**:
  - A banner that won't go away unless users agree with the terms.

Comprehension (C.x)

Threats in this category relate to the comprehension of privacy notices and choice mechanisms.

[C.1] Contradictory Statements or Implementations

[C.1.1] Conflicting Statements
Definition:

- Conflicting statements within the privacy notice and choice mechanisms about the same data practice; this includes statements that users are likely to interpret as conflicting even if they can also be interpreted in a way that is not conflicting.

Evaluation Questions:

- Do the privacy notice and choice mechanisms include statements that might be **reasonably interpreted as contradictory** by a user?

Examples:

- **Ineffective**:
    - When a company claims, "We do not collect personal data," yet still collects email addresses from users during registration.

[C.1.2] Mismatched Notice Statement and Choice Implementation
Definition:

- Statements in privacy notices regarding available choices and how they are implemented are inconsistent with the service/device's actual choice implementation.

Evaluation Questions:

- If a privacy notice indicates that there is a way for users to control certain data practices, is that **achievable in the choice interfaces**?

Examples:

- **Ineffective**:
    - Statement: “Users are able to withdraw their consent to share their personal information.”
        - Actual practice: The website offers no option for users to withdraw their consent.
    - Statement: “You do not have to consent in order to obtain any products or services.”
        - Actual practice: Users who do not click the consent button are not able to proceed with their order.
    - Statement: “You can opt-out of email communications from us in our privacy preference center.”
        - Actual practice: Some email communication choices are available in the privacy preference center but newsletter choices are provided only on the newsletter page of the website.

[C.2] Inconsistent Terminology
Definition:

- Different terms used throughout the notice and choice interface to describe the same concept or data type.

Evaluation Questions:

- Do the privacy notices or choice interfaces exhibit inconsistency by **using different terms interchangeably** for the same concept or type of data?

Examples:

- **Ineffective**:
  - A privacy notice describes a user's available "opt-out" choices regarding service personalization; however, the privacy choice mechanisms use the term "unsubscribe" to substitute "opt-out," leading to potential confusion among users.
  - A service's privacy notices and choice mechanisms refer to third-party data sharing using different names ("data sharing," "data disclosure," "data partnerships," etc.) without including an explanation in regards to the connection between these concepts.

[C.3] Difficult to Understand

[C.3.1] Unclear Terms, Statements, or Choices

Definition:

- Privacy notices and choice mechanisms containing unclear words, phrases, or statements that could lead to confusion, ambiguity, unclearness, or multiple interpretations.

Evaluation Questions:

- Are there **hedging words** in the privacy notices or choice mechanisms? (e.g., may, would, possible, could, etc.)
- Do the terms used in the notices or choice mechanisms have **multiple possible meanings** or **interpretations**?
- Do the notices and choice mechanisms make it clear **to whom they are referring** when discussing different parties involved with certain data practices, such as **third parties**?

Examples:

- **Ineffective**:
  - A privacy choice mechanism offering an opt-out for third-party data data sharing states that enabling the opt-out will ensure "unnecessary" third-party data sharing will be disabled, without clearly defining what "unnecessary" means.
  - A privacy notice for a personal fitness tracking app states that a user's location data is shared with "relevant parties" in order to improve user experience, without defining who "relevant parties" are.

[C.3.2] Use of Legal or Technical Jargon

Definition:

- Privacy notices and choice interfaces that use jargon or acronyms that could make it challenging for the intended audience to understand the content without providing informative and non-disrupting explanations. This includes both technical terms that require an expert level of knowledge and legal terms that may not be familiar to most people.

Evaluation Questions:

- Do the privacy notices and choice mechanisms contain:
    - Legal jargon or technical terms difficult for average readers to understand?
    - Many legal clauses or subclauses?
    - References to laws or regulations without explanation?
- Are terms and acronyms that are difficult for an average user to understand **accompanied by appropriate definitions and explanations**?

Examples:

- **Ineffective**:
    - Legal jargon
        - "In the event of a **force majeure** event, we shall not be liable…"
        - "**Notwithstanding** anything to the contrary **herein**, we reserve the right to retain your data…."
    - Technical jargon
        - "We collect minimal personal data and employ technologies like **encryption** and **local computing** for enhanced privacy"

[C.3.3] Use of Complex Sentences

Definition:

- Privacy notices and choice interfaces that use language that may be challenging for the intended audience to understand due to the use of long or complex sentence structures or uncommon words.

Evaluation Questions:

- Are the sentences in the notices and choice mechanisms **clear and straightforward**, or are they overly long or complex that may be difficult for users to grasp the main idea?
- What is the **reading level** of the text in the privacy notices and choice mechanisms? (You can use online reading level checkers or tools built into word processors to check this.)

Examples:

- **Ineffective**:
    - "Please be aware of the fact that should you elect to preclude us from acquiring your personal information, it may consequently preclude our ability to offer specific experiences, products, and services to you, thereby potentially compromising the personalization and efficacy of our offerings."

[C.4] Consequences not Adequately Explained

Definition:

- The consequences of privacy choice options are not clearly explained to users in their presented context.

Evaluation Questions:

- Does the system provide insights into the **likely outcomes** of each user's choice?

Examples:

- **Ineffective**:
    - The following statement appears in isolation without an accompanying policy: "By clicking the AGREE button, you agree to the collection and use of information in accordance with our policy."

[C.5] Inadequate Feedback

Definition:

- Privacy choice mechanisms provide none or insufficient feedback in terms of whether the privacy settings have been successfully updated after users submit their choices or info regarding the current state of privacy settings.

Evaluation Questions:

- Does the system/service confirm **that their privacy preferences** have been successfully updated? (e.g., popup notices, icons, emails, etc.)
- Does the system/service promptly offer **transparent and timely feedback** to users?
- Can users readily **check the current state** of their privacy settings?

Examples:

- **Effective**:
  - An IoT device has a camera and a microphone that can be turned on and off, and there is a green light on the IoT device, as well as the screen on the mobile, when the camera or the microphone is active. (as shown in the figure below).
- **Ineffective**:
  - Users can fill out a form to ask that their data be deleted from a service, but there is no confirmation when the form is processed or a settings page that indicates data has been deleted.
  - A social media platform allows users to control the audience for their content, but their current audience settings are not shown or linked in the content posting interface.
  - An IoT device has a camera that can be turned on and off, but there is no indicator of when the camera is on and recording (as shown in the figure below).

[C.6] Confusing Buttons/Toggles/Checkbox
Definition:

- Choice mechanisms that are presented in a confusing way resulting in user uncertainty as to which state represents each choice.

Evaluation Questions:

- Are the choice mechanisms implemented effectively so that **users can tell if a choice option indicates "yes" or "no" by the design** (color, style, text labels, etc.)?
- Can users **easily understand** what will be selected once the button, toggles, or checkbox has been set to a particular value?

Examples:

- **Ineffective**:
  - When the toggle is not labeled with words, or it doesn't match the corresponding text description (e.g., when the toggle is switched to "on" but it means opt-out instead of opt-in, as shown in the figure below), it can be difficult to determine their state based only on position or color.
  - A cookie banner with an unlabelled X or a close button that does not convey to users what choice is made when they close the banner.

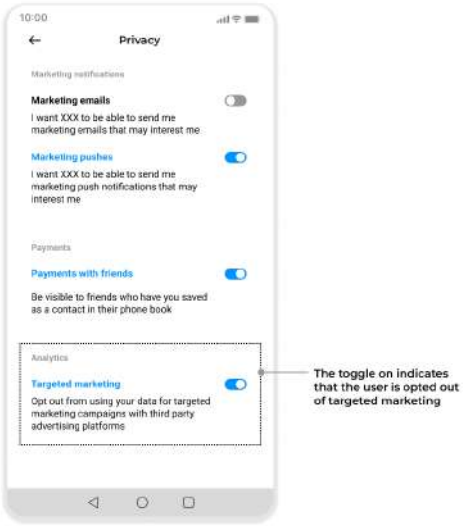

Appropriate Choices (AC.x)

Threats in this category are related to lack of appropriate privacy choice mechanisms that both comply with legal regulations and match user expectations.

[AC.1] Limited Choice
Definition:

- Privacy choice mechanisms that lack or fail to adequately cover privacy choice options that are required by applicable law or expected by users.

Evaluation Question(s):

- Does the system/service offer users choices with regard to **relevant data practices or processes** based on the context?
- Does the system/service fail to provide mechanisms for users to express their preferences regarding data practices for which choices are required per applicable laws or where choices would be reasonably expected by users?
    - Data collection, storage, and/or processing
    - Third-party data sharing and/or selling
    - Data deletion
    - Cookie and tracking mechanisms

Examples:

- **Ineffective**:
    - No choice is provided with regards to key aspects of users' concerns such as data sharing with 3rd parties.
    - For AI systems, users should be given the choice to opt-in/opt-out of being included in the training data set.

[AC.2] Excessive or Redundant Choice Options

Definition:

- Privacy choice mechanisms provide too many choices or require too much effort for users to make effective decisions or exercise certain privacy rights.

Evaluation Questions:

- Does the system/service overwhelm users with **an excessive number of privacy choices** without a reasonable approach to simplify them (such as clicking a certain button to select a bundle of choice options), potentially impeding their decision-making process?
- Does the system/service force users to **fill in an excessive number of elements/forms/requirements,** potentially putting too much of a burden on users (e.g., deletion requests)?

Examples:

- **Ineffective (excessive):**
    - An online shopping platform provides users with the following privacy options:
        - "Allow personalized product recommendations."
        - "Enable suggestions based on your browsing history."
        - "Receive tailored product suggestions."

      which essentially provide the same functionality.

[AC.3] Inadequate or Excessive Granularity

Definition:

- Privacy choice options that either fail to encompass user expectation of choice (inadequate granularity) or present too many fine-grained choices (excessive granularity), rendering it unsatisfying or confusing for users when making choices.

Evaluation Questions:

- Are the available options for users **overly extreme** and with no middle ground, therefore not capable of aligning with users' needs?
- Are users **exclusively offered two distinct options** without any middle ground or customization possibilities? (e.g., Accept/Decline, Yes/No)
    - **Note:** If the situation logically necessitates only two possible options and allows users to communicate their privacy preferences effectively using binary choices, this threat can be ignored.
- Are the available options for users **excessively detailed and nuanced**, providing a range of choices that may not be aligned with users' needs?

Examples:

- **Ineffective (inadequate):**
    - In an IoT environment, for instance, primary users are sometimes presented with **two extreme choices**: either allow their guests to use their accounts with full access or have them use guest accounts that have strict restrictions regarding the functions they can use.
    - For location-based services, having the option to **share location only while actively using** the application rather than allowing the application to track location constantly.
- **Ineffective (excessive):**
    - A choice interface requires users to allow or reject trackers from **50 different third-party tracking companies** if they want to opt-out of third-party tracking. (A better interface would allow users to opt-out of all tracking and/or opt-out of a small number of categories of tracking (e.g. advertising, site analytics, social media, etc.).

[AC.4] Difficult to Modify Previous Choices

Definition:

- Privacy choice mechanisms that make it difficult or impossible for users to modify their choices after submitting the choice to the system.

Evaluation Questions:

- Does the system/service **prevent the user from modifying or retracting** a privacy preference after submission?
- How long does it take an average user to find the locations to initiate choice modification, and can the average user find it at all?

Examples:

- **Ineffective**:
    - The privacy policy will only be displayed when the users use the service for the first time. Users are unable to check the privacy policy again following the same path or modify their privacy choices on the same page (as shown in the figure below).

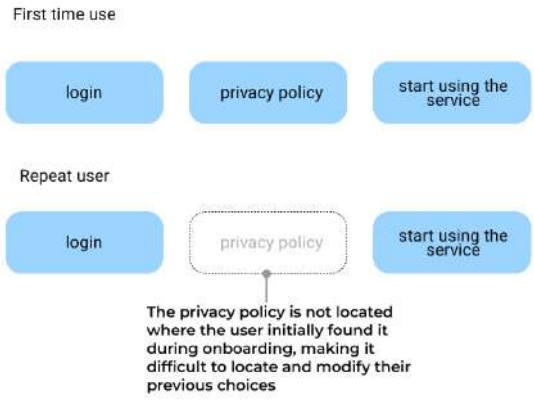

Manipulative Elements (M.x)

Threats in this category are related to manipulative interfaces in privacy notices and choice interfaces.

[M.1] Manipulative Statements
Definition:

- Privacy notices and choice interfaces that use subtle language to manipulate users into taking less privacy-protective actions.

Evaluation Questions:

- Do the privacy notices and choice mechanisms **manipulatively associate less privacy-protective actions with positive outcomes,** such as improved user experience, benefits for other users or society, or other desirable results?
- Do the privacy notices and choice mechanisms **manipulatively associate more privacy-protective actions with less positive outcomes**, such as poor user experience or missing out on benefits?
- Are there any usage of **guilt-based, manipulative language or content** in privacy choice mechanisms designed to potentially evoke negative emotions as to influence users toward taking a less privacy-protective action?

Examples:

- **Ineffective**:
    - Instead of saying "share your data," a nudged version might be phrased as "enhance your experience by sharing your data."
    - "... if you opt-out, you'll still see ads, but they won't be tailored to your interests based on your activity."

[M.2] Visually Manipulative Design
Definition:

- A deceptive/dark pattern (*dark patterns: trick users into taking an action that is not in their best interest*) where the interface encourages users to take invasive privacy actions by using

particularly enticing or noticeable font or button colors, different font or button sizes, or manipulative bundling and layouts.

Evaluation Questions:

- Does the visual representation subtly **encourage users**, particularly average users, to **select the less privacy-protective option**?

Examples:

- **Ineffective**:
  - The "Accept" button under the cookie notice is emphasized with blue color, encouraging users to click on the "Accept" button to agree to the use of cookies, instead of manually setting up their cookie preferences (as shown in the figure below).

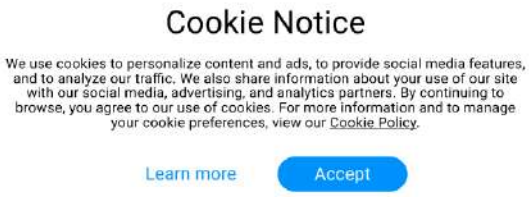

[M.3] Asymmetric Efforts Required for Different Privacy Protection Levels

Definition:

- A deceptive/dark pattern (*dark patterns: trick users into taking an action that is not in their best interest*) in which users need to take more steps for more privacy-protective actions than for less privacy-protective actions.

Evaluation Questions:

- Does taking the privacy-protective action (e.g., opt-ins, rejecting cookies) take the same amount of effort/steps as taking the privacy-invasive action (e.g., opt-outs, accepting all the cookies)?

Examples:

- **Ineffective**:
  - The user can simply click on "Accept" to agree to all the use of cookies, but needs to go down a path with multiple steps ("Learn more" -> "Cookie Preferences" -> the cookie preferences page) to modify their cookie preferences (as shown in the figure below).

[M.4] Less Privacy Protective Defaults

Definition:

- Privacy settings default to options with lower level of protections on privacy.

Evaluation Questions:

- Does the system or service offer default privacy settings that are **less privacy-protective** than other options, requiring users to adjust them for higher levels of protection manually?

Examples:

- **Ineffective**:
    - For first-time registered users, data sharing with third parties is turned on by default.

[M.5] Unexpected Choice Alteration

Definition:

- User choices that lead to unexpected consequences, especially with regard to other choices.

Evaluation Questions:

- Does the system inform users of the consequences of their choices, including **automatic changes to other settings**?
- Are users clearly notified about all changes to their settings, even those they did not directly select?
- Do **presets or hierarchical choice** interfaces clearly convey the choices associated with each top-level setting?

Examples:

- **Effective:**
    - Choices are either independent of one another or their connection has been clearly articulated, such as disabling third party sharing will automatically turn off sharing for all individual partners.
- **Ineffective:**
    - When a user opts into “using one’s activity to show a customized ad,” the system automatically adjusts other settings, such as “sharing personal data with third parties” without further notice.

APPENDIX H
TABLES

A. Demographic

		Participants
Age	18-24	1
	25-34	12
	35-44	6
	45-54	2
	55-64	2
	65+	0
	Prefer not to answer	2
	No Answer	1
Gender	Male	18
	Female	4
	Prefer not to answer	3
	No Answer	1
Highest Degree	Bachelor's Degree	5
	Master's Degree	15
	Doctorate (e.g., PhD, EdD)	1
	Professional Degree (e.g., MD, JD)	3
	Prefer not to answer	1
	No Answer	1
Occupation	Software Engineer	5
	Privacy Product Manager	2
	Privacy Engineer	10
	Security Engineer	2
	PhD Student/Trained as Privacy Engineer	1
	Privacy Officer/Attorney	6

TABLE IV
DEMOGRAPHIC DISTRIBUTION OF PARTICIPANTS

B. Ground Truth Threat Instances

TABLE V: Ground-truth threat instances for AccuFrame and Beyond. *"N_i" refers to the ith notice in that scenario and "C_i" refers to the ith choice in that scenario. "Minor: threat type x" indicates that if with-taxonomy participants put the same threat instance in type x, we regard it as correct placement during placement analysis.*

Threat Types	AccuFrame	Beyond
[DU.1] Nonexistent or Difficult to Locate		N1 - Page 3: It's difficult to find info related to audio data both throughout the user journey and inside the privacy policy (Minor: DU.6.2); Content related to information collection being too generic; C2 - Page 14: Users not being aware of the existence of choices (missing notice); C3 - Page 15: Users not being aware of the existence of choices (missing notice); Other - Page 7: The privacy policy documents being difficult to find (Minor: DU.6.2)
[DU.2] Ineffective Timing		
[DU.3] Ineffective Channel		Other: Mentioning of the channel issue (Minor: DU.6.2)
[DU.4] Lack of Centralized Management		Other - Notices and choices in the Nexa app with confusing names (Minor: C.3.1); mentioning of choices spread across multiple channels
[DU.5] Decoupled Notice and Choice		
[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids	N1 - Page 20: No effective navigation like a table of contents but long text	N2 - Page 10; Other - Too wordy (especially the 2 choice texts)
[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)	C1 - Page 3: Too many links, no highlighting of important info; C2 - Page 16: Too many clicks to get to this page; the arrow buttons require extra effort to learn more about cookies; C3 - Page 20: Users have to take all the steps, including reading all the data rights, to be able to access info related to their actual data rights (Minor: DU.1); should also highlight important info such that it's less time-consuming for users to access the data deletion control	N1 - Page 1: Collapse button or too many links or no link (Minor: M.1); C1 - Too many clicks/steps to get to the settings, including both before getting to the Nexa app, unclear voice command, and after arriving at the app and having to navigates through multiple pages (Minor: DU.1)
[DU.7] Poorly Formatted Notices and Choices		
[DU.8] Dysfunctional components (links, buttons, switches, etc)		
[DU.9] Distracting Visual/Audio Effects		
[C.1.1] Conflicting Statement(s)	N1 - Page 3 and 20: It is stated on page 3 that "Any VTO Input Data collected through your use of VTO Tool will be permanently destroyed within one year after the date such Data is collected." But on page 20, it is said "We do not store any of those data...", making it likely for users to be confused about whether their biometric data has been stored or not; Page 20 states "To map your facial feature to provide recommendations...," While it is said in section one that they only use this data during VTO; Page 20 states that users have the right to withdraw their consent while in reality, they cannot do that (Minor C.1.2); Page 20 mentions information not shared with the third party but the consent (page 3) clearly includes Faceview as third party	N2 - Page 10: "To review and delete the voice recordings associated with your account (including any audio resulting from a false wake), you can go to Settings ¿ Nexa Privacy in your Nexa app." This conflicts with what's said before, which is user should be able to use voice commands for deletion directly (Minor: C.1.2)
[C.1.2] Mismatched Notice Statement and Choice Implementation		
[C.2] Inconsistent Terminology		
[C.3.1] Unclear Terms/Statements/Choice Implementation	C1 - Page 3: Not clear how privacy policies mentioned differ from one another; "which may potentially be considered of a biometric nature"; unclear about consequence; confusing terms; what the links lead to; N1 - Page 20	Other - Page 3, 10, 13, 14, 15
[C.3.2] Use of Legal or Technical Jargon	C1 - Page 3 (for all jargon, if participants mention wanting more definitions put it here instead of C.3.1 unclear terms); C2 - Page 16: technical jargon; N1 - Page 20	Other - Page 10
[C.3.3] Use of Complex Sentences	Other - Page 3, 20	
[C.4] Consequences not adequately explained		
[C.5] Inadequate Feedback		Other - No pop-up to indicate that user choice has been applied

Threat Types	AccuFrame	Beyond
[C.6] Confusing Buttons/Toggles/Checkbox	C1 - Page 3: Not sure what the "x" button means (Minor AC.1, AC.3)	N2 - Page 8: QR code (Minor: M.2). Unclear if users can still view the policy on the TV or do they have to scan the QR code; C3 - Page 15: Confusing UserName and toggle
[AC.1] Limited Choice	C1 - Missing privacy Choice related to control of the processing (use) of biometric data being collected (which includes data deletion)	N1 - Page 4: No opt out of audio data collection/ choose not to be recorded when being presented with the policy; no control over what types of data users are opting to share
[AC.2] Excessive or Redundant Choice Options		
[AC.3] Inadequate or Excessive Granularity		C2 - Page 14 (Minor: AC.1)
[AC.4] Difficult to Modify Previous Choices	C1 - Page 3 (Minor: DU.1, C1)	
[M.1] Manipulative Statements		Other - Page 3, 10, 12, 14, 15
[M.2] Visually Manipulative Design	C2 - Page 16: Toggles turned on but are gray (Minor: C.6); users have to click the arrow buttons to view more specific details	
[M.3] Asymmetric Effort required for Different Privacy Protection Levels		N1 - Page 1: Use of skip and select all button (Minor: DU.6.2, M.1, M.2)
[M.4] Less Privacy Protective Defaults	C2 - Page 16: Cookies on by default	C2 - Page 14; C3 - Page 15
[M.5] Unexpected Choice Alteration		

C. Relevant Threat Instances Identified by Participants But Not Considered Ground Truth

TABLE VI: Ground-truth threat instances for AccuFrame and Beyond identified by participants. *"N_i" refers to the ith notice in that scenario and "C_i" refers to the ith choice in that scenario. "Minor: threat type x" indicates that if with-taxonomy participants put the same threat instance in type x, we regard it as correct placement during placement analysis*.

Threat Types	AccuFrame	Beyond
[DU.1] Nonexistent or Difficult to Locate		
[DU.2] Ineffective Timing	C2: Not showing cookie banner (including when mentioning Chloe has not been actively prompted to make a cookie choice)	C1: Deletion by voice control should be presented earlier in the process
[DU.3] Ineffective Channel		
[DU.4] Lack of Centralized Management	N1 - Page 18: Multiple privacy policy pages that are confusing	
[DU.5] Decoupled Notice and Choice	C2 - Page 16: Mentioning of rights to opt-out of info sold, but can't find that option	N1 - Page 3: Mentioning of rights to access/delete recordings, but can't find that option; N2 - Page 3: Mentioning "Data privacy queries" exists, but can't find that option; can't find where and how to set up voice ID; C2 - Page 14: "This will not delete other information about sounds you choose for Nexa to detect such as your Guard preferences and device settings." It is not clear where the user made this choice and how to locate it
[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids	C1 - Page 3: Too much text/ bad way of organizing text; N1 - Page 20: No effective navigation like a table of contents but long text	N1 - Page 20: Should list out collected data types in a more clear manner
[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)		N2 - Page 10: Too many links to other policies; no link to instruction on how to delete for all marketplace
[DU.7] Poorly Formatted Notices and Choices		
[DU.8] Dysfunctional components (links, buttons, switches, etc)		
[DU.9] Distracting Visual/Audio Effects		
[C.1.1] Conflicting Statement(s)		
[C.1.2] Mismatched Notice Statement and Choice Implementation		
[C.2] Inconsistent Terminology	C1 - Page 3: "agree" vs. "expressively agree"; "facial data" vs "biometric data"	

Table VI continued from previous page

Threat Types	AccuFrame	Beyond
[C.3.1] Unclear Terms/Statements/Choice Implementation		
[C.3.2] Use of Legal or Technical Jargon		
[C.3.3] Use of Complex Sentences		
[C.4] Consequences not adequately explained	C1 - Page 3: Not clear what will happen if select "not agree" to the policies, e.g., unable to use the VTO feature; C2 - Page 16: What will happen if users reject the cookies	C2 - Page 14: The three options are unclear, e.g., "save for 12 months and delete older history"; C3 - Page 15: "If you turn this off, voice recognition and new features may not work well for you. " Unclear what "not work well" means
[C.5] Inadequate Feedback		
[C.6] Confusing Buttons/Toggles/Checkbox	C2 - Page 16: Not clear if users have to click "save preferences" button to save changes to the cookie options	N1 - Page 1 and 2: Confusing button "Next"
[AC.1] Limited Choice		C2 - Page 14: Lack of choices on deleting/not recording detected sounds
[AC.2] Excessive or Redundant Choice Options		
[AC.3] Inadequate or Excessive Granularity	C2 - Page 16: No reject all button	C1 - Page 13: Lack of granularity to allow for deletion of selected voice recordings; C2 - Page 14 (Minor: AC.1)
[AC.4] Difficult to Modify Previous Choices		
[M.1] Manipulative Statements	Other - Terms wrapped in quotation marks; "subject to ..."; page 21: Mentioning of losing rewards account	
[M.2] Visually Manipulative Design		
[M.3] Asymmetric Effort required for Different Privacy Protection Levels		N2 - Page 10: Users have to delete data associated with all marketplaces separately
[M.4] Less Privacy Protective Defaults		
[M.5] Unexpected Choice Alteration		

D. Importance Rating

Threat Type	Average Accu-With Importance Rating	Average Accu-No Importance Rating	Average Beyond-With Importance Rating	Average Beyond-No Importance Rating
[DU.1] Nonexistent or Difficult to Locate			4.30	3.83
[DU.2] Ineffective Timing				
[DU.3] Ineffective Channel			3.97	4.17
[DU.4] Lack of Centralized Management	3.73	3.10	3.67	3.67
[DU.5] Decoupled Notice and Choice				
[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids	3.33	4.00	N/A	3.00
[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)	3.51	3.69	3.71	4.42
[DU.7] Poorly Formatted Notices and Choices				
[DU.8] Dysfunctional components (links, buttons, switches, etc)				
[DU.9] Distracting Visual/Audio Effects				
[C.1.1] Conflicting Statement(s)	5.00	4.00	4.50	4.00
[C.1.2] Mismatched Notice Statement and Choice Implementation				
[C.2] Inconsistent Terminology				
[C.3.1] Unclear Terms/ Statements	3.81	4.13	4.39	3.67
[C.3.2] Use of Legal or Technical Jargon	3.55	3.81	3.50	4.17
[C.3.3] Use of Complex Sentences	3.67	N/A		
[C.4*] Consequences not adequately explained				
[C.5*] Inadequate Feedback			3.38	4.00
[C.6*] Confusing Buttons/Toggles/Checkbox	3.58	N/A	3.00	N/A
[AC.1*] Limited Choice	4.33	4.00		
[AC.2*] Excessive or Redundant Choice Options				
[AC.3*] Inadequate or Excessive Granularity			2.00	2.00
[AC.4*] Difficult to Modify Previous Choices	4.37	4.75		
[M.1] Manipulative Statements			3.00	N/A
[M.2*] Visually Manipulative Design	3.78	3.50		
[M.3*] Asymmetric Effort required for Different Privacy Protection Levels			4.00	5.00
[M.4*] Less Privacy Protective Defaults	3.80	3.67	3.25	4.00
[M.5*] Unexpected Choice Alteration				

TABLE VII

AVERAGE IMPORTANCE RATING FOR EACH GROUND-TRUTH THREAT INSTANCE BY PARTICIPANTS FROM THE FOUR CONDITION GROUPS. AN N/A INDICATES THAT EITHER NO PARTICIPANTS FROM THAT GROUP IDENTIFIED ANY GROUND TRUTH THREAT INSTANCES RELATED TO THAT THREAT TYPE OR PARTICIPANTS WHO MENTIONED IT DID NOT PROVIDE AN IMPORTANCE RATING. A BLANK CELL MEANS THE THREAT TYPE IS NOT APPLICABLE TO A PARTICULAR SCENARIO.

E. Recall for Four Threat Categories

Participant Number	DU Recall	C Recall	AC Recall	M Recall
AW1	80.00%	25.00%	100.00%	50.00%
AW2	60.00%	62.50%	100.00%	100.00%
AW3	60.00%	62.50%	50.00%	50.00%
AW5	40.00%	37.50%	50.00%	50.00%
AW6	100.00%	100.00%	50.00%	100.00%
AW7	60.00%	25.00%	50.00%	50.00%
AN1	40.00%	12.50%	0.00%	100.00%
AN2	20.00%	12.50%	50.00%	0.00%
AN3	40.00%	37.50%	0.00%	0.00%
AN4	80.00%	0.00%	0.00%	0.00%
AN5	20.00%	25.00%	100.00%	50.00%
AN6	20.00%	50.00%	0.00%	0.00%
AN7	80.00%	25.00%	0.00%	100.00%
BW1	50.00%	66.67%	50.00%	100.00%
BW3	40.00%	50.00%	50.00%	75.00%
BW4	50.00%	50.00%	0.00%	25.00%
BW5	60.00%	33.33%	50.00%	50.00%
BW7	30.00%	33.33%	50.00%	25.00%
BW8	40.00%	66.67%	0.00%	25.00%
BN1	60.00%	50.00%	0.00%	0.00%
BN2	20.00%	33.33%	50.00%	25.00%
BN3	40.00%	33.33%	50.00%	50.00%
BN4	30.00%	0.00%	50.00%	25.00%
BN5	30.00%	16.67%	0.00%	50.00%
BN6	20.00%	50.00%	100.00%	25.00%
BN7	40.00%	33.33%	0.00%	0.00%

TABLE VIII
RECALL FOR FOUR THREAT CATEGORIES BY EACH PARTICIPANT IN THE FOUR CONDITION GROUPS

	DU Recall	C Recall	AC Recall	M Recall
AW Average	66.67%	52.08%	66.67%	66.67%
AN Average	42.86%	23.21%	21.43%	35.71%
BW Average	45.00%	50.00%	33.33%	50.00%
BN Average	34.29%	30.95%	35.71%	25.00%

TABLE IX
AVERAGE RECALL FOR FOUR THREAT CATEGORIES BY THE FOUR CONDITION GROUPS

F. Percentage of Participants Identifying Each Threat Type

AccuFrame - Threat Types	Average percentage of participants from Accu-With who identified threat instances in this threat type	Average percentage of participants from Accu-No who identified threat instances in this threat type	Beyond - Threat Types	Average percentage of participants from Beyond-With who identified threat instances in this threat type	Average percentage of participants from Beyond-No who identified threat instances in this threat type
[DU.1] Nonexistent or Difficult to Locate	N/A	N/A	[DU.1] Nonexistent or Difficult to Locate	54.17%	39.29%
[DU.2] Ineffective Timing	N/A	N/A	[DU.2] Ineffective Timing	N/A	N/A
[DU.3] Ineffective Channel	N/A	N/A	[DU.3] Ineffective Channel	83.33%	42.86%
[DU.4] Lack of Centralized Management	83.33%	71.43%	[DU.4] Lack of Centralized Management	50.00%	42.86%
[DU.5] Decoupled Notice and Choice	N/A	N/A	[DU.5] Decoupled Notice and Choice	N/A	N/A
[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids	50.00%	14.29%	[DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids	0.00%	14.29%
[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)	66.67%	42.86%	[DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy)	50.00%	35.71%
[DU.7] Poorly Formatted Notices and Choices	N/A	N/A	[DU.7] Poorly Formatted Notices and Choices	N/A	N/A
[DU.8] Dysfunctional components (links, buttons, switches, etc)	N/A	N/A	[DU.8] Dysfunctional components (links, buttons, switches, etc)	N/A	N/A
[DU.9] Distracting Visual/Audio Effects	N/A	N/A	[DU.9] Distracting Visual/Audio Effects	N/A	N/A
[C.1.1] Conflicting Statement(s)	50.00%	28.57%	[C.1.1] Conflicting Statement(s)	66.67%	42.86%
[C.1.2] Mismatched Notice Statement and Choice Implementation	N/A	N/A	[C.1.2] Mismatched Notice Statement and Choice Implementation	N/A	N/A
[C.2] Inconsistent Terminology	N/A	N/A	[C.2] Inconsistent Terminology	N/A	N/A
[C.3.1] Unclear Terms/ Statements	41.67%	42.86%	[C.3.1] Unclear Terms/ Statements	100.00%	71.43%
[C.3.2] Use of Legal or Technical Jargon	61.11%	23.81%	[C.3.2] Use of Legal or Technical Jargon	16.67%	57.14%
[C.3.3] Use of Complex Sentences	50.00%	0.00%	[C.3.3] Use of Complex Sentences	N/A	N/A
[C.4*] Consequences not adequately explained	N/A	N/A	[C.4*] Consequences not adequately explained	N/A	N/A
[C.5*] Inadequate Feedback	N/A	N/A	[C.5*] Inadequate Feedback	66.67%	14.29%
[C.6*] Confusing Buttons/Toggles/Checkbox	50.00%	0.00%	[C.6*] Confusing Buttons/Toggles/Checkbox	25.00%	0.00%
[AC.1*] Limited Choice	50.00%	14.29%	[AC.1*] Limited Choice	33.33%	57.14%
[AC.2*] Excessive or Redundant Choice Options	N/A	N/A	[AC.2*] Excessive or Redundant Choice Options	N/A	N/A
[AC.3*] Inadequate or Excessive Granularity	N/A	N/A	[AC.3*] Inadequate or Excessive Granularity	33.33%	14.29%
[AC.4*] Difficult to Modify Previous Choices	83.33%	28.57%	[AC.4*] Difficult to Modify Previous Choices	N/A	N/A
[M.1] Manipulative Statements	N/A	N/A	[M.1] Manipulative Statements	66.67%	0.00%
[M.2*] Visually Manipulative Design	50.00%	28.57%	[M.2*] Visually Manipulative Design	N/A	N/A
[M.3*] Asymmetric Effort required for Different Privacy Protection Levels	N/A	N/A	[M.3*] Asymmetric Effort required for Different Privacy Protection Levels	33.33%	14.29%
[M.4*] Less Privacy Protective Defaults	83.33%	42.86%	[M.4*] Less Privacy Protective Defaults	50.00%	42.86%
[M.5*] Unexpected Choice Alteration	N/A	N/A	[M.5*] Unexpected Choice Alteration	N/A	N/A

TABLE X
PERCENTAGE OF PARTICIPANTS IDENTIFYING EACH THREAT TYPE FOR EACH SCENARIO

G. Irrelevancy Analysis

Code Name	Definition
Not privacy related	Participants discussed issues that was not privacy-related
Outside of storyboard	Participants wanted some mechanisms that happened outside of the user journey (e.g., user login process)
Not specific to the notices and choices	Participants mentioned some issues that were not relevant to the specific notice or choice that we asked them to focus on
Misunderstand the storyboard	Participants misinterpreted certain components in the storyboard (e.g., ignoring the existence of a button, not understanding the user persona's action)
Missing disclosure	Participants wanted some information related to data practices (e.g., data retention, security safeguard, data controller's contact information, back-end mechanisms, legal compliance) that did not exist in the notice
Data practice criticism	Participants disagreed with the data practices
Not a threat (disagree with participants)	The two coders unanimously disagreed with the participant that a threat instance brought up was really a threat
Detailed definitions for "Not a threat"	The participant claimed something was missing or unclear, but we disagree; the suggestion would reduce privacy or usability; the proposal is unworkable due to a system misunderstanding; some may prefer this, but it's not clearly better; no impact on privacy or usability; the suggestion is unclear

TABLE XI
IRRELEVANCY CODEBOOK

Participant Number	Total Number of Irrelevant Threats Identified	Not Privacy Related	Outside of Storyboard	Not Specific to the Notices and Choices	Misunderstand the Storyboard	Missing Disclosure	Data Practice Criticism	Not a Threat (Disagree with Participants)
AW1	4	0	0	1	1	0	0	2
AW2	3	0	0	0	0	1	1	1
AW3	7	0	0	1	1	0	1	4
AW5	0	0	0	0	0	0	0	0
AW6	3	0	0	1	0	0	1	1
AW7	2	0	0	0	1	1	0	0
AN1	2	0	0	0	0	2	0	0
AN2	9	0	1	3	0	3	2	0
AN3	2	0	0	0	0	0	0	2
AN4	8	0	1	1	0	2	2	2
AN5	5	0	0	2	0	1	1	1
AN6	2	0	0	0	0	1	0	1
AN7	6	0	0	0	0	2	1	3
BW1	6	2	0	0	2	1	1	0
BW3	8	0	1	1	3	0	1	2
BW4	1	1	0	0	0	0	0	0
BW5	4	1	0	0	1	0	0	2
BW7	4	0	0	0	2	0	1	1
BW8	0	0	0	0	0	0	0	0
BN1	1	0	0	0	0	0	0	1
BN2	6	0	0	0	0	5	0	1
BN3	1	0	0	0	0	0	0	1
BN4	4	0	0	2	0	0	1	1
BN5	5	1	0	1	1	1	0	1
BN6	12	1	0	3	0	1	6	1
BN7	11	3	1	3	0	4	0	0

TABLE XII
NUMBER OF IRRELEVANT ISSUES IDENTIFIED BY PARTICIPANTS

H. Codebooks for Participants' Threat Identification Experience

Overall Category	Definition
Framework	*(Only applies to with-taxonomy groups)* Discussion regarding the use of the framework taxonomy
Approach	*(Only applies to no-taxonomy groups)* Approach used during the threat identification process
Experience	*(Only applies to no-taxonomy groups)* Experiences during the threat identification process (e.g., easy, hard, familiarity of the threat identification process)
Procedure	Comments/thoughts on the interview procedure, e.g., too many tables, too little time, not wanting to switch between tabs, wanting to get the taxonomy beforehand (for with-taxonomy participants)
Scenario	Comments/thoughts regarding the complexity or helpfulness of information included in the scenario storyboards
Suggestion	Concrete solutions and actionable suggestions

TABLE XIII: Codebook for threat identification experience category

Higher-level Category	Code Name	Code Code Definition
Approach	approach_framework	Participants followed follow certain types of frameworks such as STRIDE, LINDDUN, or NIST
	approach_user	Participants put themselves into the shoes of the users
	approach_privacyanalyst	Participants analyzed the scenario from the perspective of privacy analysts
	approach_different_aspects	Participants analyzes the scenario from multiple (greater than 2) different perspectives (e.g., users, legal, business)
	approach_data_practices	Participants' approach focused on whether they can get a full and clear picture of the data practices involved
	approach_specific_elements	Participants tried to detect if the given scenario performed well according to following aspects: user consent, effectively informing the user, usability, contradiction between privacy policy and actual implementation, understandable language, manipulative statements, matching user expectation, potential misuse case, and fulfilling legal compliance
Experience	experience_similar	Participants have done similar tasks in the past, e.g., at work or during a course or had experience in interacting with similar apps
	experience_easy	Participants considered the process to be easy or straightforward, not mentioning any particular difficulties
	experience_moderate	Participants considered the process to be moderate, as they provided both easy and hard parts of the exercise
	experience_not_familiar_with_-framework	Participants were not familiar with the framework they brought with them
	experience_deficiency_framework	Participants felt that the framework they brought with them were too simple and incapable of capturing some of the threats
Scenario	scenario_clear	Scenario or information presented in the scenario was clear and helped participants focus
	scenario_long	The scenario was a bit long
	scenario_complicated	The scenario was complicated
	scenario_insufficient	Participants felt that the scenario was vague or insufficient to realistically convey user-oriented threats accurately
	scenario_enough	Participants believed that they have enough things to perform the analysis
Suggestion	suggestion_framework	Participants considered having a taxonomy/a more useful taxonomy to be helpful
	suggestion_backend_dataflow	Participants wanted to know how things worked on the backend, e.g., had a information flow or DFD
	suggestion_physical	Participants wanted to have an in-person session, so that they could engage in a whiteboard session or interact with physical artifacts, e.g., a toolkit or card, when performing the analysis
	suggestion_different_expert	Participants suggested having experts from different fields (e.g., legal expert, UX Designers, product managers, engineering) to be involved in the process
	suggestion_interactivity	Participants wanted to interact with the static components (i.e., links, buttons in the scenario)

TABLE XIV: Codebook for threat identification experience - no-taxonomy

Higher-level Category	Code Name	Code Definition
Framework	Framework_not_applicable	Participants were not sure how this analysis would actually apply in the real world
	Framework_application	Participants considered this taxonomy to be more applicable than other existing frameworks for real world cases
	Framework_checklist	Participants considered the taxonomy to work as a checklist (or served to help them double check if they have missed any threats
	Framework_guide	Participants considered the taxonomy to work as a guide
	Framework_detailed	Participant considered the taxonomy to be very detailed and well explained
	Framework_specific	Participants liked that the taxonomy gave really specific types of threats, like distinguishing between the different types of "difficult to understand"
	Framework_did_not_change	Participants were already aware of the threats listed in the taxonomy, so it did not change their approach on threat identification
	Framework_easy	The taxonomy was easy and intuitive to use (e.g., clear and straight-forward, easy to understand)
	Framework_help_find_more	The taxonomy helped participants consider threats that they were not aware of otherwise
	Framework_heavy	Framework taxonomy felt heavy to use (e.g., disjoint categories, too many threats in the taxonomy, too long
	Framework_not_useful	Participants explicitly said that the taxonomy was not useful (in general)
	Framework_familiar	The taxonomy became more useful after getting familiar with it
	Framework_similar	Participants have done/were doing similar work in the same/related area
	Framework_useful	Participants considered the taxonomy to be useful in general (without addressing specific reasons) or was more useful compared to other frameworks
	Framework_systematic	Participants considered the taxonomy to provide a systematic (methodical, comprehensive) way to approach threat identification
	Framework_limited_content	Participants considered the content of the taxonomy to be limited, and didn't cover all types of threats in user-oriented privacy field (e.g., specific to privacy notice only)
	Framework_limited_scope	Participants considered the current taxonomy to be only usable for specific audience
	Framework_no_consequence	Participants felt the taxonomy was not presented with the actual threats/harms (e.g., illegal for the company to do that, emotional harm for users, financial consequence) or the taxonomy needed to be provided with the level of risk for each threat
	Framework_overlapping	The definition of threats tended to overlap, resulting in the same threat instance be mapped to multiple threat categories/not sure which category to map to
	Framework_unclear_purpose	Participants were not sure about the goal/purpose of this taxonomy
Procedure	Procedure_inappropriate_format	Participants said the taxonomy could be better presented in a Figma or PPT format instead of using the table format
	Procedure_order	Participants preferred to identify the threat instance first and then mapped it to the threat types (the current interview order was weird)
	Procedure_tables	It was difficult for participants to use multiple tables for threat identification (e.g., the table was too long to read)
	Procedure_time	The time duration for the interview was too tight to complete all the tasks
	Procedure_unfamiliar_framework	Participants indicated that if they had been more familiar with the taxonomy they could have performed better in the task
	Procedure_allow_for_focus	By examining each notice and choice, participants were able to dig deeper with regard to each data practice
Suggestion	Suggestion_AI_categorization	Participants suggested that it would be easier to use the help of some tools such as AI to help with determining which taxonomy category/type a threat belongs to
	Suggestion_interview_order	A more natural order is to focus on the scenario first and then use the taxonomy to identify threats
	Suggestion_interview_tabular	Participants suggested to combine multiple tables into one
	Suggestion_interview_framework_beforehand	Participants stated wanting to get familiarize with the taxonomy before the start of the interview
	Suggestion_interview_threat_importance	Participants wanted more definitions regarding how to label threat instance's importance
	Suggestion_framework_broader_aspects	Participants suggested that the taxonomy should cover boarder aspects (business, legal, data practices, non privacy experts, designers, etc.) and can be applied to various types of audiences, instead of limiting to the user aspects
	Suggestion_framework_condense	Participants suggested condensing the threat categories and types when possible
	Suggestion_framework_consequence	Participants suggested including more details about the consequences for each threat (e.g., physical safety, emotional harms, risk if organizations fail to address a threat in the taxonomy, including legal and reputational harm)
	Suggestion_framework_flexibility	Participants considered the taxonomy could have had different versions depending on the intended user

TABLE XV: Codebook for threat identification experience - with-taxonomy