Are You Onboard? Exploring Experiential Privacy Onboarding in Virtual Reality

Shady Mansour
shady.mansour@ifi.lmu.de
LMU Munich
Munich, Germany

Verena Winterhalter
verena.winterhalter@ifi.lmu.de
LMU Munich
Munich, Germany

Florian Alt
florian.alt@ifi.lmu.de
LMU Munich
Munich, Germany
University of the Bundeswehr
Munich, Germany

Viktorija Paneva
viktorija.paneva@ifi.lmu.de
LMU Munich
Munich, Germany

Abstract

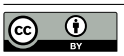
VR applications increasingly rely on sensors such as hand tracking, eye tracking, and microphones, yet users are often asked to grant sensor access before experiencing how these permissions affect interaction. This paper examines privacy onboarding as a pre-consent intervention that helps users reason about the consequences of permissions in VR. We developed an interactive onboarding prototype for a productivity scenario in which users previewed the functional effects of permission decisions before encountering standard VR permission dialogs. In a between-subjects study (N=40), we found that onboarding increased participants' perceived ability to continue using the system after refusing permissions, highlighting the importance of making refusal a visible and supported path. Interview findings further showed that onboarding shifted consent reasoning from inference based on familiar permission conventions toward demonstrated trade-offs between data access and functionality. While onboarding did not significantly affect task completion time, self-reported understanding, or overall permission-granting rates, the findings suggest that privacy onboarding can support more deliberate consent reasoning by making permission consequences more experiential. We discuss implications for designing VR consent experiences that support autonomy, consequence understanding, and ongoing consent management.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; *Mixed / augmented reality*; *Virtual reality*; • **Security and privacy** → **Usability in security and privacy**; *Privacy protections.*

Keywords

Virtual Reality, Privacy, Onboarding, Consent, Privacy Permissions, Usable Privacy



This work is licensed under a Creative Commons Attribution 4.0 International License.
VRST '26, Sendai, Japan
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2811-2/2026/11
https://doi.org/10.1145/3822517.3848675

ACM Reference Format:
Shady Mansour, Verena Winterhalter, Florian Alt, and Viktorija Paneva. 2026. Are You Onboard? Exploring Experiential Privacy Onboarding in Virtual Reality. In *32nd ACM Symposium on Virtual Reality Software and Technology (VRST '26), November 16–18, 2026, Sendai, Japan.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3822517.3848675

1 Introduction

Virtual reality (VR) systems increasingly rely on sensors that are close to the body, continuous, and difficult for users to inspect. Head and hand motion can identify individuals at scale [21, 34]; gaze and body tracking can expose sensitive behavioral signals [6]; and avatar gaze can even leak typed input in attacks such as *GAZEploit* [49]. These capabilities make permission decisions consequential, but the moment of consent is often compressed into short, front-loaded prompts that ask users to decide before they have experienced the system.

This mismatch is especially problematic in immersive environments. Privacy notices and consent forms already fail to attract sustained attention on conventional screens [30, 36], and users often struggle to connect permissions with later application behavior [40, 44]. In VR, the challenge is not only whether a prompt is readable, but whether it helps users understand how abstract data access changes embodied interaction. A microphone permission may determine whether voice input is available; hand tracking may change how objects can be manipulated; eye tracking may enable gaze-based interaction or analytics. Users, therefore, need to reason about both privacy risk and functional consequence.

Prior work has argued that privacy in extended reality (XR) should be treated as a process that unfolds across onboarding, permission prompts, in-experience controls, and post-session review [29]. This paper focuses on the first stage of that process: privacy onboarding. Drawing on VR onboarding research [13], the central design question is whether a system can prepare users for permission decisions by showing the consequences of consent before the prompt appears. We introduce *experiential privacy onboarding*, a pre-consent approach in which users experience the interaction consequences of granting or denying permissions through interactive demonstrations before making a consent decision.

We developed an interactive onboarding prototype for a VR productivity scenario using three common permission types: microphone, hand tracking, and eye tracking. The *Onboarding* condition presented short visual demonstrations and optional explanations showing what each permission enabled and what would remain possible if access was denied. A *Baseline* condition used the same VR-style prompts without the experiential onboarding. This design isolates the contribution of onboarding while keeping the consent interface close to contemporary all-at-once permission practice.

A between-subjects user study with 40 participants examined three research questions:

RQ1: Does experiential privacy onboarding increase users' perceived control and voluntariness when granting or denying VR permissions?

RQ2: Does experiential privacy onboarding affect users' perceived understanding and reasoning about the consequences of permission decisions?

RQ3: Does experiential privacy onboarding influence permission-granting behavior?

The findings suggest that privacy onboarding primarily affected how participants reasoned about permission decisions rather than whether they granted permissions. Participants in the *Onboarding* condition reported significantly greater perceived ability to continue using the system after refusing permissions. Interviews further showed that *Baseline* participants often inferred consequences from familiar permission patterns, whereas *Onboarding* participants reasoned from demonstrated tradeoffs between functionality and data access. While the onboarding did not significantly affect task completion time and self-reported understanding, these findings highlight the value of making refusal visible as a supported path.

Contribution Statement: The contribution of this paper is threefold: **(1)** the introduction and implementation of *experiential privacy onboarding*, a pre-consent approach that demonstrates permission consequences before consent; **(2)** a controlled user study comparing this onboarding against a standard permission flow; and **(3)** design implications for VR permission systems that emphasize supported refusal, consequence-based understanding, and accountability.

2 Related Work

Privacy in XR can be understood not as a one-off consent event but as a continuous, multi-stage process unfolding throughout use [25, 31]. Building on this view, Mansour et al. propose a User Privacy Journey for XR privacy that spans onboarding, permission prompts, in-experience controls, and post-experience review [29]. This paper focuses on the first stage, privacy onboarding, and reviews work that directly motivates this focus: XR privacy risks, the limits of static consent, and experiential approaches to onboarding and privacy education.

2.1 Privacy Risks and the Limits of Static Consent in XR

XR systems embed dense sensor arrays close to the body, continuously capturing head and hand motion, eye gaze, and audio that can be re-purposed far beyond their immediate use. Systematizations of VR privacy show that short windows of routine telemetry expose a wide range of inferable attributes [32]. Motion data can identify users at a large scale [21, 34]; eye and body tracking can reveal sensitive behavioral signals [6]; and avatar gaze can leak typed input [49]. Yet audits of commercial VR ecosystems find that app privacy policies rarely disclose the sensor data collected or the inferences drawn from it [33, 48]. Users are therefore asked to consent to data practices they cannot see and whose downstream implications are opaque.

Conventional consent mechanisms are poorly suited to this problem. Users often ignore text policies and lengthy consent forms [30, 36], and even well-studied notice-design principles remain difficult to satisfy on 2D screens [22, 42]. In XR, text-heavy consent becomes still less appropriate because users are immersed, perceptually occupied, and often learning unfamiliar interactions [24]. Hence, we treat consent as an ongoing process rather than a single transaction, and examine whether the *onboarding* that precedes consent can improve the quality of subsequent permission decisions.

2.2 Permission Models and Immersive Consent

Permission systems vary in granularity, from coarse bundles to fine-grained per-sensor controls, and in timing, from all-at-once (AAO) requests at first launch to just-in-time (JIT) prompts shown when a feature needs access [14, 15]. On mobile platforms, contextual and developer-explained requests can improve attention and align permissions with user expectations [40, 44, 45, 50], whereas front-loaded requests risk habitual acceptance and privacy fatigue [43]. Immersive systems add challenges: users may be steered by deceptive or manipulative spatial designs [19, 26], and design practitioners report normalizing such tactics under business pressures [51].

This study does not compare permission timing or prompt design. Instead, both conditions use identical AAO consent prompts modeled on contemporary VR operating systems. This allows the study to isolate a narrower question: whether experiential onboarding before the prompt changes how users understand and reason about the same permission request.

2.3 Onboarding and Privacy Education in VR

Onboarding teaches users how to interact with an unfamiliar system and is especially consequential in VR, where users are isolated from the surrounding world and must learn new input modalities. Chauvergne et al. [13] synthesize VR onboarding practices into design dimensions including context, interactivity, flexibility, assistance, coverage, feedback, input, and information medium. Prior systems show the value of embedded, multimodal instruction for VR tasks, including video-based tutorials augmented with in-context 3D aids [46, 47] and comparisons of instruction modalities for immersive interaction [11]. This work informs the prototype design, but it primarily addresses task onboarding rather than privacy.

Privacy-specific onboarding in VR remains scarce. Lim et al. [27] show that role-playing tutorials can raise privacy awareness by letting users experience how their data could be misused. Security-education research similarly suggests that video- and demonstration-based delivery can motivate protective behavior more effectively than text [4, 5]. Most directly related to this paper, Abraham et al. [1] let users experience how different permission levels affect both

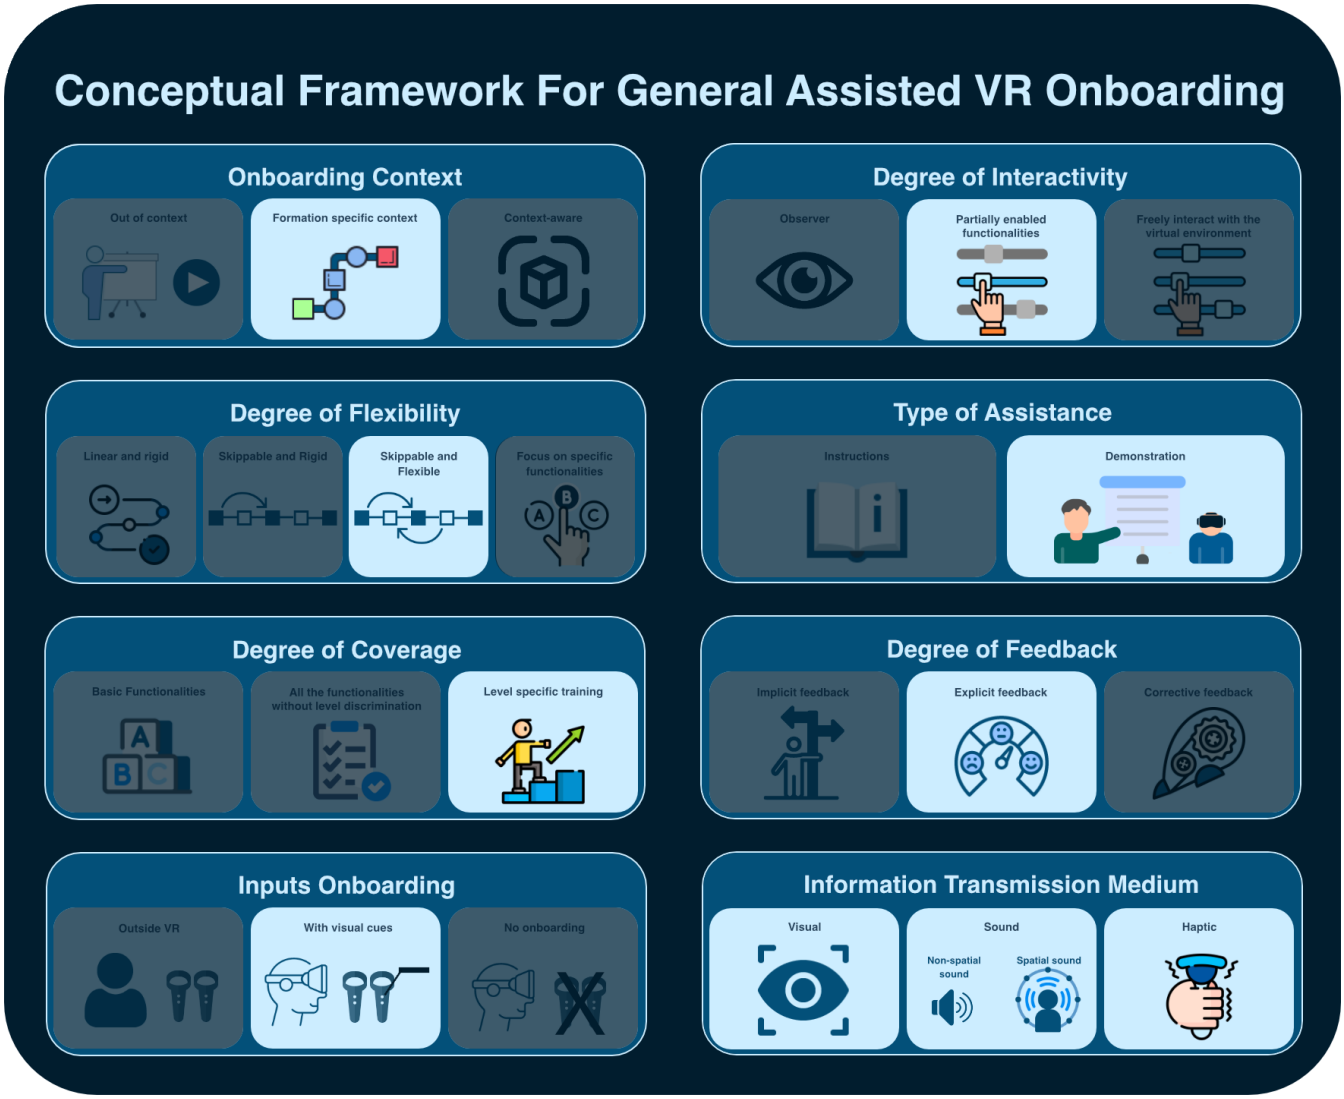


Figure 1: Conceptual framework for general assisted VR onboarding, adapted from [13]. The framework identifies eight design dimensions: *Onboarding Context*, *Degree of Interactivity*, *Degree of Flexibility*, *Type of Assistance*, *Degree of Coverage*, *Degree of Feedback*, *Inputs Onboarding*, and *Information Transmission Medium*. Highlighted dimensions indicate the design choices implemented in our privacy onboarding prototype.

privacy and functionality in everyday XR, finding such experiential permission interfaces more informative and trustworthy than smartphone-style prompts. We extend this direction to the onboarding stage, before the permission prompt appears.

The expected effect of onboarding also depends on users' mental models of VR data. Prior studies report limited awareness of what XR sensors capture and infer [3, 20], and users often reason by analogy to smartphone permissions even though continuous, multimodal sensing can support richer behavioral and biometric profiling [38, 39, 41]. Users may also overestimate their privacy knowledge, making subjective confidence a poor proxy for comprehension [2, 35]. These findings motivate evaluating not only whether onboarding makes users feel informed, but whether it changes the basis of their consent reasoning. The study therefore asks whether interactive, demonstration-based onboarding improves perceived control (RQ1), understanding of consent consequences (RQ2), and permission-granting behavior (RQ3).

3 Designing Experiential Privacy Onboarding

We build on the User Privacy Journey by Mansour et al. [29], which translates their broader Multi-Layered Privacy Permission Framework into a sequence of user-facing privacy encounters. The journey begins with *onboarding and education*, where users first encounter what data a system may request and why. It then moves to

contextual permission prompts, where users make concrete access decisions; *in-experience transparency and control*, where systems communicate ongoing data use and allow adjustments during use; and *post-experience review*, where users can revisit, audit, or revise earlier decisions. We focus on the first stage, as onboarding shapes the expectations users bring into every later consent decision.

In this work, we operationalize *experiential privacy onboarding* as a pre-consent intervention that makes the consequences of permission decisions concrete and observable through interactive demonstrations within the target environment, allowing users to preview both what each permission enables and what remains possible without it before encountering the permission prompt. We implement this approach in a VR productivity scenario with three common permission types: microphone input, hand tracking, and eye tracking. Unlike conventional notices that describe permissions abstractly, our onboarding lets users preview permission outcomes as interaction consequences inside the VR environment.

We based the design of our privacy onboarding process on the VR onboarding framework by Chauvergne et al. [13]. Figure 1 maps our prototype to the eight design dimensions of the framework. The onboarding uses a *formation-specific context*: a dedicated Privacy Onboarding Lab within the VR application for learning and exploration. It offers *partially enabled functionalities*: participants can interact with the station objects, while controller-based teleportation between predefined locations restricts locomotion to preserve experimental control. The onboarding is *skippable and flexible*, allowing participants to explore the voice, hand-tracking, and eye-tracking stations in any order, revisit them, or proceed directly to the exit. It uses *demonstrations* as its form of assistance, allowing users to preview the consequences of granting and denying each permission rather than only reading about them. The design provides *level-specific coverage*: short demonstrations convey the core interaction consequences to all participants, while optional *Read More* panels provide additional information about data use, collection, retention, risks, and user control. The demonstrations provide *explicit feedback* by showing the permission-enabled interaction and the available fallback interaction. Visual cues within the demonstrations introduce the permission prompts and their *Allow* and *Don't Allow* actions. Information is communicated through visual, auditory, and haptic channels to support multimodal learning in VR. Together, these choices adapt the framework to pre-consent privacy onboarding by making both permission states observable before a consent decision is required.

The design was guided by four goals. First, onboarding should make permission consequences *experiential* rather than merely descriptive. Second, it should preserve *autonomy*: declining a permission should remain a usable path, not a hidden failure state. Third, it should support *layered disclosure*: all users receive short demonstrations, while users who want more detail can access additional information on demand. Fourth, it should avoid collecting the very data it explains before consent has been requested. The onboarding environment therefore used only controller-based interaction until participants reached the actual permission prompts.

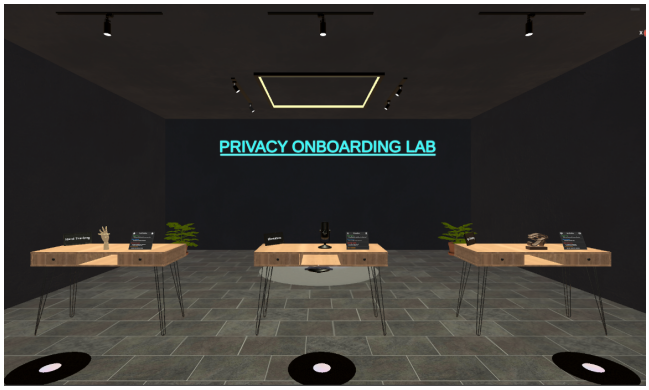


Figure 2: The Privacy Onboarding Lab environment showing three privacy-related information stations with interactive objects representing hand tracking (left), microphone (center), and eye tracking (right).

3.1 Grounding the Permission Scenario

To make the prototype resemble plausible VR use rather than an artificial privacy exercise, we grounded the permission scenario in a productivity context. We inspected ten productivity applications from the Meta Quest Store that were also available on at least one additional VR or AR platform. Following the design space for privacy choices by Feng et al. [16], the inspection recorded the requested permissions, the timing of requests, the type of choice offered, and whether the app explained why data access was needed.

We used this scan as design grounding rather than as a representative market audit. Across the inspected applications, microphone, hand-tracking, and eye-tracking access appeared repeatedly in productivity-oriented interaction, while permission requests were typically binary and front-loaded at or near first launch. Explanations of how the data would affect functionality were limited. Microphone input, hand tracking, and eye tracking were therefore selected as the prototype's permission set, and the consent flow was modeled as separate all-at-once prompts before the main task environment. This allowed us to isolate the onboarding intervention while keeping the permission experience close to contemporary VR platform practice.

3.2 The Privacy Onboarding Lab

The *Privacy Onboarding Lab* was the entry point for participants in the *Onboarding* condition (Figure 2). It presented three stations corresponding to the three requested sensing modalities: hand tracking, microphone input, and eye tracking. Each station used a physical metaphor for the data type, such as a hand model, microphone, or eye symbol, so that the permission being explained was spatially and visually anchored in the environment.

Participants navigated the lab via predefined teleportation points, a common locomotion technique to reduce discomfort in VR [10]. They could visit the stations in any order, revisit them, or skip onboarding by moving directly to the exit.

At each station, grabbing the corresponding object triggered a short demonstration. The room lighting dimmed, the active station

was highlighted, and a video played on a holographic stage. Each demonstration used the same structure:

(1) a preview of what would happen if the user selected *Allow*;
(2) a demonstration of the resulting interaction benefit, such as voice dictation or hand gestures;
(3) a preview of what happens upon selecting *Don't Allow*;
(4) a demonstration of the fallback interaction available without that permission.

We used in-world videos because permission consequences often involve dynamic changes in interaction that are difficult to communicate through static text or images. Videos also allowed us to present permission-enabled and fallback interactions in a standardized manner without activating the corresponding sensors before users have provided consent.

The key design choice was to demonstrate both permission states symmetrically. The onboarding did not present denial as a failure or as a blocked path; instead, it showed which interactions would be unavailable and which alternatives would remain. To support layered disclosure, each station also included an optional *Read More* panel with additional information about the purpose of data use, what data would and would not be collected, possible risks, data retention, and user control.

The lab itself did not collect consent and did not activate the microphone, hand-tracking, or eye-tracking input. Consent was separated from onboarding to ensure that both experimental conditions encountered identical prompts at the same point in the study. This avoided a privacy contradiction in which the system would collect sensitive input before users made the permission decision.

A walkthrough of the onboarding is available in the supplementary video.

3.3 Consent Prompts

After onboarding, participants encountered three VR-style permission prompts requesting access to microphone input, hand tracking, and eye tracking. *Baseline* participants encountered the same prompts without first entering the onboarding lab. Each prompt used a uniform layout modeled on contemporary VR operating-system permission dialogs: a modality-specific title, a short explanation of the requested access, and binary *Allow* and *Don't Allow* actions. The prompts were identical across conditions. Screenshots of all consent prompts are provided in the supplementary material.

3.4 Simulated Productivity Tasks in VR

The post-consent task environment was a simulated VR office (Figure 3). We chose a productivity setting because it plausibly motivates microphone, hand-tracking, and eye-tracking permissions without relying on entertainment-specific mechanics. We designed the tasks so that granting a permission could improve or diversify interaction, but denying it would not prevent task completion.

In the first task, participants created and organized virtual sticky notes about personal productivity. They could enter note text via a virtual keyboard or, if microphone access was granted, via voice dictation. In the second task, participants evaluated the notes by marking agreement or disagreement. Depending on their permission choices, they could use controller buttons, hand gestures (such as thumbs-up and thumbs-down), or gaze-based selection.



Figure 3: Simulated productivity environment in VR. Participants completed two tasks: creating and organizing reflective notes, and evaluating notes using the available interaction modalities.

This task design made the consequences of permission observable during use. Allowing a permission enabled additional interaction modalities; denying it left controller-based alternatives available. The prototype, therefore, operationalized a central design principle for autonomy-preserving consent in VR: permissions should enhance the experience without making refusal equivalent to exclusion from the system.

4 Study Method

Our between-subjects study evaluated whether experiential privacy onboarding changes how users understand and act on VR permission requests.

4.1 Study Design

The study compared two conditions. In the *Baseline* condition, participants encountered the VR permission prompts and then completed the productivity tasks. In the *Onboarding* condition, participants first explored the Privacy Onboarding Lab, then encountered the same prompts, and then completed the same tasks. The study used a mixed-methods design combining permission decisions, interaction logs, self-report measures, and post-study interviews.

Participants were allocated to conditions in alternating order by participant number, resulting in 20 participants per condition.

4.2 Participants

We recruited 40 participants (22 females, 18 males) through university mailing lists and communication channels. Participants in the *Baseline* condition had a mean age of 26.70 years ($SD = 6.43$); participants in the *Onboarding* condition had a mean age of 27.40 years ($SD = 6.61$). Both groups had comparable prior VR experience: in each condition, 16 participants reported prior VR use, and 4 reported no prior VR experience. Participants were compensated according to the local standard of participant reimbursement. The study was approved by the institution's IRB.

The Affinity for Technology Interaction (ATI) scale [17] and the Internet Users' Information Privacy Concerns (IUIPC) scale [28] were additionally collected to characterize the sample. The groups

Table 1: Results for self-report measures by condition.

Measure	Condition	Mean	SD	Min.	Max.
Age	Baseline	26.70	6.43	18.00	43.00
	Onboarding	27.40	6.61	19.00	43.00
ATI_total	Baseline	3.98	0.86	1.89	4.89
	Onboarding	3.84	0.80	2.22	5.56
IUIPC_CTRL	Baseline	6.05	0.72	4.33	7.00
	Onboarding	5.75	1.08	4.00	7.00
IUIPC_AWA	Baseline	6.15	0.78	4.00	7.00
	Onboarding	6.27	0.75	4.67	7.00
IUIPC_COLL	Baseline	5.39	1.26	2.75	7.00
	Onboarding	5.55	1.46	2.50	7.00
IUIPC_total	Baseline	5.82	0.78	4.10	6.90
	Onboarding	5.83	0.86	3.60	7.00

were similar on ATI (*Baseline*: $M = 3.98$, $SD = 0.86$; *Onboarding*: $M = 3.84$, $SD = 0.80$) and overall IUIPC (*Baseline*: $M = 5.82$, $SD = 0.78$; *Onboarding*: $M = 5.83$, $SD = 0.86$), suggesting comparable technology affinity and baseline privacy concern across conditions.

4.3 Apparatus and Implementation

We implemented the VR environments in Unity using C# [18]. Sessions used a Meta Quest Pro standalone headset with dual QLED displays, a 90Hz refresh rate, and a per-eye resolution of 1800×1920 pixels. The headset supports the sensing modalities used in the prototype, including hand tracking and eye tracking. Participants used controller input before consent and, after the permission prompts, could use the modalities enabled by their own permission decisions.

4.4 Procedure

Participants first completed a demographic questionnaire, the ATI scale, and questions about prior VR experience. On arrival, they provided written informed consent for the study and received a short introduction to the task sequence and basic VR controls. To reduce demand characteristics [37], we told participants that the study examined interaction patterns in VR productivity applications; the privacy focus was disclosed after the interview.

Inside VR, participants completed eye-tracking calibration and then entered the experimental flow. Participants in the *Onboarding* condition first explored the Privacy Onboarding Lab (subsection 3.2), then encountered the consent prompts (subsection 3.3). Participants in the *Baseline* condition encountered the same prompts without onboarding. Both groups then completed the simulated productivity tasks (subsection 3.4).

After the VR tasks, all participants completed post-experience questionnaires about subjective understanding of consent effects and voluntariness, followed by the IUIPC scale. Participants in the *Onboarding* condition additionally completed an onboarding evaluation questionnaire and the System Usability Scale (SUS) [12]. Each session ended with a semi-structured interview about participants' reasoning, perceptions of control, and concerns about the consent process; the interview guides are available in the supplementary materials. Sessions averaged 38.65 minutes in the *Baseline* condition (range: 25–60) and 44.5 minutes in the *Onboarding* condition (range: 33–60).

4.5 Measures

We collected three types of data. First, the system logged behavioral measures: permission decisions, task completion times, use of permission-enabled features, and onboarding engagement events such as whether participants interacted with videos or optional information panels. Second, self-report measures captured participants' perceived understanding and autonomy. Two study-specific 5-point Likert item sets were used: Subjective Understanding of Consent Effects (SUCE; five items) and Voluntariness (VOL; three items). For the *Onboarding* condition, a custom onboarding evaluation covering clarity and information quality, ease of interaction, and confidence and agency was administered, along with SUS for usability. The supplementary material provides the complete questionnaire items. Third, interviews provided qualitative evidence about participants' mental models, decision strategies, and perceptions of the onboarding experience. We treated task completion time as an exploratory measure of downstream interaction performance rather than hypothesizing a directional effect: onboarding could reduce uncertainty about available interactions, but it could also encourage more deliberate behavior.

4.6 Analysis

We analyzed quantitative data in R. Because Likert responses were ordinal and several measures violated normality according to Shapiro-Wilk tests, Mann-Whitney U tests were used for between-condition comparisons. For significant effects, rank-biserial correlation is reported [23], interpreting $|r| \geq 0.1$, 0.3, and 0.5 as small, medium, and large effects. We treat the custom SUCE and VOL items as exploratory measures of perceived understanding and autonomy rather than as validated scales.

We analyzed the post-study interviews using thematic analysis [9]. Three researchers independently open-coded 20% of the interviews and discussed the codes to construct an initial codebook. The researchers then applied and iteratively revised this codebook while coding the remaining interviews. The team reconvened to consolidate overlapping codes, group them into higher-level themes, resolve disagreements, and verify representative excerpts against participants' accounts. The complete final codebook is available in the supplementary material.

4.7 Limitations

Several constraints shape the interpretation of the findings. With 20 participants per condition, the study had limited power to detect small effects or support subgroup analyses, including comparisons based on engagement with optional onboarding content. The between-subjects design reduced carryover, but individual differences may still have affected outcomes despite similar age, ATI, IUIPC, and VR-experience distributions. Because we assessed participants' mental models only after the VR experience, we could not determine how their understanding changed over the course of the study. Finally, participants were recruited primarily from a university context, limiting generalizability to broader populations,

including children, older adults, and users with different cultural or institutional trust backgrounds.

The laboratory setting also matters. Participants explicitly consented to participate in the research, and several interview comments indicated that an academic context made data sharing feel safer than it would in a commercial setting. Real-world deployments with unfamiliar providers, monetization pressures, and longer-term use may produce different permission behavior.

The prototype captured a single application genre and a single session. Results may differ in social VR, entertainment, health, fitness, or enterprise contexts, where the perceived value and sensitivity of permissions may change.

Finally, one implementation constraint affected the permission architecture: enabling microphone input replaced keyboard input rather than adding voice as an additional option. This frustrated some participants and may have amplified the importance of clearly explaining fallback behavior. The issue is treated as a limitation, but it also revealed an important expectation: users often assume permissions add functionality rather than replacing existing interaction paths.

5 Results

We organize the results around the three research questions. First, quantitative measures show that onboarding did not significantly change task completion time, self-reported understanding, or overall permission-granting rates, but did significantly increase participants' perceived ability to continue using the system after refusing permissions. Second, behavioral logs show high permission-granting rates in both conditions and uneven engagement with optional onboarding materials. Third, interview data explain the mechanism behind these results: onboarding shifted consent reasoning from inference about what might happen to demonstrated tradeoffs between functionality and data access. We provide detailed descriptive statistics, questionnaire-item tables, and distribution plots in the supplementary materials.

5.1 Behavioral Outcomes and Engagement

Task completion time did not differ significantly between conditions. Mann–Whitney tests found no significant condition differences for Task 1 (*Baseline*: $M = 376.8$ s, $SD = 159.4$; *Onboarding*: $M = 322.2$ s, $SD = 108.7$; $U = 245.0$, $p = .231$) or Task 2 (*Baseline*: $M = 234.1$ s, $SD = 103.6$; *Onboarding*: $M = 189.2$ s, $SD = 34.2$; $U = 246.0$, $p = .221$). The results do not indicate that onboarding impaired subsequent task performance.

Permission granting was high in both conditions, suggesting that onboarding did not produce a simple increase or decrease in willingness to share data. In the *Baseline* condition, 17 participants allowed hand tracking, 17 allowed eye tracking, and 16 allowed microphone access. In the *Onboarding* condition, 18 participants allowed hand tracking, 18 allowed eye tracking, and 17 allowed microphone access.

Within the *Onboarding* condition, engagement was higher for short demonstrations than for optional text panels. Most participants interacted with at least one video, and 15 of 20 interacted with all three videos. By contrast, 10 of 20 participants did not interact with any "Read More" panel; because panel interaction only indicates that the panel was held, these counts are an upper bound on attention to detailed text. Detailed engagement logs are provided in the supplementary materials. The pattern supports a core design implication of the prototype: demonstration-based onboarding was more visible and more likely to be engaged with than passive textual disclosure.

5.2 Onboarding Usability and Perceived Clarity

Participants in the *Onboarding* condition rated the prototype as usable and clear, though not frictionless. The mean SUS score was 70.25 ($SD = 13.33$), which falls within the acceptable range [8, 12] and near the lower bound of "Good" adjective ratings [7]. Participants rated the onboarding highly for explaining which data would be collected (CIQ–1: $M = 4.65$, $SD = 0.75$) and how consent decisions would affect the VR experience (CIQ–2: $M = 4.55$, $SD = 0.76$). Ratings were more moderate for confidence in making data-sharing decisions (CA–1: $M = 3.75$, $SD = 1.07$) and intuitive navigation (EI: $M = 3.60$, $SD = 1.23$), indicating the concept was understandable but the interaction design still left room for refinement.

5.3 Subjective Understanding of Consent Effects

Self-reported understanding was high in both conditions. Mean ratings were slightly higher in the *Onboarding* condition for several items, such as understanding why data would be needed (*Baseline*: $M = 3.80$, $SD = 0.89$; *Onboarding*: $M = 3.95$, $SD = 0.95$; $U = 177.0$, $p = .521$) and which features would be available depending on their choices (*Baseline*: $M = 4.15$, $SD = 0.99$; *Onboarding*: $M = 4.50$, $SD = 0.69$; $U = 165.0$, $p = .301$). However, none of the SUCE items differed significantly between conditions. We note that the SUCE items measure perceived understanding, not whether participants could accurately predict system behavior or explain data risks.

5.4 Voluntariness

Voluntariness ratings showed the clearest quantitative effect (Figure 4). Two items exhibited ceiling effects and did not differ significantly: whether granting or denying each permission was entirely the participant's choice (*Baseline*: $M = 4.80$, $SD = 0.41$; *Onboarding*: $M = 4.95$, $SD = 0.22$; $U = 170.0$, $p = .164$), and whether participants felt pressured to accept permissions (*Baseline*: $M = 4.80$, $SD = 0.52$; *Onboarding*: $M = 4.85$, $SD = 0.49$; $U = 190.5$, $p = .672$). However, participants in the *Onboarding* condition rated the statement that they could continue using the system after refusing specific permissions significantly higher than participants in the *Baseline* condition (*Baseline*: $M = 4.15$, $SD = 0.75$; *Onboarding*: $M = 4.70$, $SD = 0.80$; $U = 109.0$, $p = .005$, $r = 0.46$). This suggests that onboarding primarily affected perceived functional autonomy: participants better understood that refusal was a viable path.

5.5 Semi-structured Interviews

A thematic analysis of the interviews contextualized the quantitative results. Rather than showing that onboarding made users uniformly more privacy-protective or uniformly more willing to grant permissions, the interviews suggest that onboarding changed the basis on which participants reasoned about consent. Four themes

Table 2: Voluntariness measures and Mann-Whitney U test results by condition. $^*p < 0.05$.

	Condition	Mean	SD	Min.	Max.	Mann–Whitney U	p
VOL–1	Baseline	4.80	0.41	4.00	5.00	170.0	.164
	Onboarding	4.95	0.22	4.00	5.00		
VOL–2	Baseline	4.80	0.52	3.00	5.00	190.5	.672
	Onboarding	4.85	0.49	3.00	5.00		
VOL–3*	Baseline	4.15	0.75	3.00	5.00	109.0	**.005**
	Onboarding	4.70	0.80	2.00	5.00		

Note. $^*p < .05$ (two-tailed, Mann–Whitney U test).

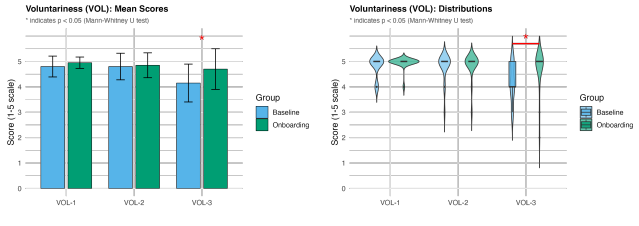


Figure 4: VOL measures by condition. VOL–3 increased significantly in the *Onboarding* condition ($^*p < 0.05$), while VOL–1 and VOL–2 exhibited ceiling effects.

are most relevant to the paper's contribution: understanding consequences, mental models of VR data, privacy-functionality tradeoffs, and revisiting consent over time.

Understanding Consequences and Impact of Consent Choices. *Baseline* participants often described the permission prompts as insufficient for anticipating consequences. One participant stated that the consent form "*doesn't provide much detail about the consequences on their experience*" (P19). When *Baseline* participants reported understanding the prompts, that understanding often came from prior technical knowledge, experimenter explanations, or personal inference rather than from the prompt itself. For example, one participant attributed awareness to being "*a computer scientist, so in the technical field*" (P43).

Onboarding participants more often grounded their understanding in what had been shown. Participants reported that onboarding clarified "*which kind of information will be collected and how [their] choices affect what will be in the future*" (P18), and described the videos as useful because "*everything was very visual*" (P38). This does not mean that comprehension was complete; some participants still wanted more detail about the extent of the collection. However, the source of understanding differed: *Baseline* participants inferred consequences, whereas *Onboarding* participants referred to demonstrated consequences.

Mental Models of VR Data and Privacy. Across both conditions, participants often treated VR-specific data as less sensitive than conventional personal information. Several participants described eye tracking, hand tracking, or microphone recordings as data that "*do not involve [their] personal information*" (P30). A small number articulated more specific risks, such as concern that audio data could be used for voice simulation because "*AI can simulate people's voice*" (P48), but such threat models were not common. This pattern suggests that onboarding about functionality did not, by itself, overcome broader gaps in VR privacy literacy.

Prior expectations also shaped consent behavior. Several participants expected permissions to be mandatory for the application to work. One *Baseline* participant explained that they "*thought that for eye tracking, if [they] didn't allow it the app wouldn't work*" (P25). This expectation matters because it turns a formally optional prompt into a perceived requirement.

Engagement and Confidence. *Baseline* participants described familiar patterns of low attention to consent information. Several acknowledged skipping consent text to continue with the experience, including one participant who described being "*someone who sometimes skips over reading to get to the good part*" (P27). In the *Onboarding* condition, participants instead described the physical and visual demonstration as engaging. One participant reported that the object-based interaction "*created a connection between the object [representing the data type] and what [they were] seeing [in the video]*" (P36), while another said the videos showed "*what specific functions [they] could use or could not use when [they] agree or disagree*" (P32).

The same demonstrations also affected confidence. *Baseline* participants tended to ground confidence in external trust signals, such as whether a company promised not to sell data or whether settings could be changed later. *Onboarding* participants additionally grounded confidence in demonstrated fallback behavior. One participant reported that confidence increased after seeing that denying data access still allowed continued use of the app (P4). Another valued that the onboarding showed "*both scenarios very equally*" (P42). This helps explain the significant VOL–3 result: the onboarding made refusal legible as a supported option.

Privacy-Functionality Tradeoffs in Permission Decisions. Both groups negotiated a privacy-functionality tradeoff, but the onboarding changed how that tradeoff was evaluated. Some *Baseline* participants described permission granting as a forced choice; one participant said their "*hand is forced to give consent*" (P47). Some granted permissions preemptively, including one participant who allowed all permissions in case they "*eventually would need it*" (P29). Participants also acknowledged that the research context made them more permissive than they might be in a commercial context.

Onboarding participants also recognized the tradeoff, but some described it through demonstrated benefits and alternatives. One participant explained that the onboarding "*showed an advantage*" of allowing permissions (P18), while another said that visualizing the benefit led to a different choice than usual (P22). Although VOL–2 showed no between-condition difference in reported pressure and exhibited a ceiling effect, some participants described fallback demonstrations as a source of autonomy. For example, one participant said the onboarding "*let [them] think for [themselves]*" because it showed that the system remained usable without the permission (P42). This pattern suggests that onboarding did not simply increase privacy concern or increase disclosure; it made permission decisions more deliberate.

Revisiting and Managing Consent Over Time. Participants in both conditions recognized that permission choices may need to be revisited. Common triggers included discovering that the app was

not usable with the current settings, learning more about data practices, changes to terms or context, and simply forgetting the consequences of earlier decisions. One participant noted that after enough time they might forget what had been accepted and need to revisit the choice (P1).

The *Onboarding* condition added a further trigger: expectation verification. Participants described wanting to revisit onboarding or settings when the application behaved differently from what had been demonstrated. One participant would return "*when [they] feel that something is weird in the app*" (P4). Another suggested returning when new functionality was introduced (P12). Hence, the onboarding functioned not only as pre-consent instruction but also as a reference point against which later system behavior could be judged.

Participants' improvement suggestions reflected this difference. *Baseline* participants requested simpler text, visual explanations, a later revision, and support for deciding later. *Onboarding* participants focused less on replacing the approach and more on refining it: making the interactive objects more noticeable, using clearer headings for video segments, replacing videos with 3D animation, and allowing users to test interactions directly. Several also requested verification mechanisms, such as the ability to audit which data had actually been collected. These suggestions point beyond onboarding toward accountable permission systems.

6 Discussion

The findings suggest that experiential privacy onboarding is most valuable as an autonomy- and sense-making intervention, rather than as a mechanism that simply increases privacy concern or reduces permission granting. Participants who experienced onboarding did not report significantly higher understanding or grant fewer permissions. However, they were significantly more likely to perceive refusal as compatible with continued system use, and interview findings indicate that their decisions were grounded in demonstrated consequences rather than inferences drawn from brief permission prompts. In this section, we discuss four implications of these findings: the distinction between perceived and experiential understanding, the role of functional autonomy in consent, the limits of layered disclosure, and the need for verification mechanisms beyond onboarding.

6.1 Perceived Understanding Is Not Enough

The absence of significant SUCE differences is not evidence that onboarding had no effect. Rather, it exposes a measurement problem for consent research in VR: users can feel that they understand a permission even when the basis of that understanding differs substantially. *Baseline* participants often inferred consequences from familiar permission patterns, whereas *Onboarding* participants referred to specific demonstrations. Both forms of reasoning can produce high self-report scores, but they are not equivalent.

This distinction matters because VR permissions connect abstract data access to embodied interaction. A text prompt can state that hand tracking enables gesture input, but an experiential preview can show what becomes easier, what remains possible without the permission, and where fallback interactions are available. These findings indicate that perceived understanding alone may be an insufficient evaluation metric for privacy interfaces, as it does not necessarily reflect users' ability to make informed privacy decisions. Future evaluations should include behavioral comprehension checks, such as asking users to predict what will happen under different permission states, rather than relying only on confidence or agreement items.

Recommendation 1

Evaluate privacy onboarding by assessing whether users can explain or predict permission consequences, rather than relying on perceived understanding or confidence alone.

6.2 Functional Autonomy Is Central to Meaningful Consent

The strongest quantitative result was the VOL–3 effect: participants in the *Onboarding* condition were significantly more likely to perceive refusal as compatible with continued system use. This finding points to a core problem in consent design: **presenting two options does not necessarily create a meaningful choice**. If users believe that refusal will break the application, the choice is only nominal. The onboarding reduced this functional opacity by showing both the permission-enabled interaction and the fallback interaction.

This has a direct design implication for immersive systems: consent interfaces should communicate **not only what data access enables, but also what remains available without it**. Denial should be designed and presented as a supported path. This is especially important in VR, where users may be reluctant to interrupt the experience, may have limited ability to inspect settings, and may assume that sensor access is necessary for basic functionality. Autonomy-preserving permission systems should thus keep core functionality available and make that fallback visible before the decision is made.

Recommendation 2

Show both permission-enabled interactions and usable fallback interactions before requesting consent.

6.3 Layered Disclosure Still Has Attention Limits

The onboarding design combined short demonstrations with optional "Read More" panels, allowing participants to access additional information. While this layered structure helped to avoid overwhelming users by placing all information into the main flow, the interaction logs suggest that passive availability alone is not enough. Half of the participants in the *Onboarding* condition did not interact with any optional panel. Separately, seven participants (P04, P10, P20, P22, P30, P36, and P48) explicitly expressed a need or desire for additional information about permissions or data practices. Of these, four interacted with a Read More panel (P04, P10, P22, and P30), whereas three did not (P20, P36, and P48). This indicates that some participants needed additional information but did not access the optional written panels, illustrating that participants did not always seek out additional information even when

they recognized a need for it. Video engagement, by contrast, was considerably higher, with 90% of participants interacting with at least one video.

Together, these condition-level patterns highlight a persistent challenge in privacy communication: making additional information available does not ensure that users will seek it during onboarding. For VR privacy onboarding, this suggests that textual detail should remain available as a supporting resource, but not as the primary mechanism for communicating decision-relevant consequences. Additional information requires stronger discoverability, better legibility, and more appropriate timing within the user journey. Some privacy information may also become more relevant after initial use, when users have a concrete reason to revisit a permission. Privacy onboarding should therefore be treated as the first touchpoint in an ongoing consent process rather than as the single moment where all privacy information must be communicated and understood, aligning with prior work that conceptualizes consent as a continuous and evolving process rather than a one-time decision [25, 29, 31].

Recommendation 3

Use demonstrations for decision-relevant consequences and reserve optional text for supporting detail.

6.4 Functional Understanding Does Not Guarantee Privacy Literacy

The onboarding helped participants reason about the functional consequences of permission decisions, but it did not fully address the sensitivity of the underlying VR sensor data. Participants in both conditions often described hand tracking, gaze, and voice as interaction inputs rather than data streams that may support biometric identification, inference, or profiling. This observation aligns with prior work showing a gap between the technical sensitivity of VR telemetry and users' intuitive privacy assessments [21, 33].

Recommendation 4

Explain privacy implications in addition to functional effects, including the inferences and profiling that sensor data may support.

Participants also wanted evidence that systems honor permission decisions. Requests for audit logs, sensor-use indicators, and mechanisms to verify which data has been stored suggest that trust cannot be established through explanation alone. While demonstrations can make permission consequences more tangible, participants' comments point to a complementary need for accountability mechanisms that allow users to inspect and verify data practices over time.

Recommendation 5

Connect onboarding to later indicators, controls, and review mechanisms that let users inspect whether permission decisions are honored.

6.5 Future Work

The findings motivate several directions for future research. First, future studies could examine privacy onboarding as part of a longer consent journey rather than as a single pre-consent intervention. Longitudinal and in-the-wild deployments could investigate whether experiential understanding persists over time, how users revisit or revise permission decisions, and how onboarding interacts with later privacy touchpoints such as contextual reminders, in-experience controls, and post-session review.

Second, future work could develop more reliable measures of consent quality. The present study suggests that users can report high levels of understanding while relying on very different forms of reasoning. Future evaluations could therefore move beyond self-report and assess whether users can accurately predict system behavior, explain permission consequences, and understand the implications of alternative permission states.

Finally, future work could investigate how to communicate the inferential sensitivity of VR data without overloading users. The present study suggests that demonstrations are effective for explaining functionality, but additional design work could explore how to explain biometric and behavioral inference in ways that are concrete, timely, and actionable.

7 Conclusion

This paper examined privacy onboarding as a pre-consent intervention for VR permission decisions. Our interactive onboarding prototype demonstrated the consequences of granting or denying microphone, hand-tracking, and eye-tracking permissions before participants encountered standard VR-style permission prompts.

In a between-subjects study with 40 participants, we observed no statistically significant differences in task completion time, self-reported understanding, or permission-granting rates. The clearest effect was on perceived autonomy: participants who experienced onboarding were more likely to understand that the system remained usable after refusing permissions. Qualitative findings further revealed that onboarding changed the basis of consent reasoning. Whereas participants in the *Baseline* condition often inferred consequences from familiar permission conventions, participants in the *Onboarding* condition reasoned about demonstrated trade-offs between functionality and data access.

The findings suggest that the added benefit of experiential privacy onboarding lies not in changing what users choose, but in supporting how they choose. For VR designers, the implication is direct: permission systems should make both acceptance and refusal visible, experiential, and supported. Meaningful consent in immersive systems depends not only on explaining what data access enables, but also on helping users understand what remains possible when access is denied.

Acknowledgments

The authors would like to thank all participants in the user study. This research received funding from the German Research Foundation (DFG) under grant agreement no. 521584224.

References

[1] Melvin Abraham, Mark Mcgill, and Mohamed Khamis. 2024. What You Experience is What We Collect: User Experience Based Fine-Grained Permissions for Everyday Augmented Reality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 772, 24 pages. doi:10.1145/3613904.3642668

[2] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514. arXiv:https://www.science.org/doi/pdf/10.1126/science.aaa1465 doi:10.1126/science.aaa1465

[3] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 427–442. https://www.usenix.org/conference/soups2018/presentation/adams

[4] Elham Al Qahtani, Lipsarani Sahoo, Yousra Javed, and Mohamed Shehab. 2022. "Why would Someone Hack Me out of Thousands of Students": Video Presenter's Impact on Motivating Users to Adopt 2FA. In *Proceedings of the 27th ACM on Symposium on Access Control Models and Technologies* (New York, NY, USA) *(SACMAT '22)*. Association for Computing Machinery, New York, NY, USA, 139–150. doi:10.1145/3532105.3535013

[5] Yusuf Albayram, John Liu, and Stivi Cangonj. 2021. Comparing the Effectiveness of Text-based and Video-based Delivery in Motivating Users to Adopt a Password Manager. In *Proceedings of the 2021 European Symposium on Usable Security* (Karlsruhe, Germany) *(EuroUSEC '21)*. Association for Computing Machinery, New York, NY, USA, 89–104. doi:10.1145/3481357.3481519

[6] Samantha Aziz and Oleg Komogortsev. 2025. Exploring the Uncoordinated Privacy Protections of Eye Tracking and VR Motion Data for Unauthorized User Identification. arXiv:2411.12766 [cs.HC] https://arxiv.org/abs/2411.12766

[7] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: Adding an adjective rating scale. *Journal of usability studies* 4, 3 (2009), 114–123.

[8] Aaron Bangor, P Kortum, JA Miller, et al. 2008. The system usability scale (SUS): An empirical evaluation. *International Journal of Human-Computer Interaction* 24, 6 (2008), 574–594.

[9] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI research: Going behind the scenes*. Morgan & Claypool Publishers, San Rafael, CA, USA.

[10] Evren Bozgeyikli, Andrew Raij, Srinivas Katkoori, and Rajiv Dubey. 2016. Point & Teleport Locomotion Technique for Virtual Reality. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play* (Austin, Texas, USA) *(CHI PLAY '16)*. Association for Computing Machinery, New York, NY, USA, 205–216. doi:10.1145/2967934.2968105

[11] Lal Bozgeyikli, Andrew Raij, Srinivas Katkoori, and Redwan Alqasemi. 2017. Effects of Instruction Methods on User Experience in Virtual Reality Serious Games. In *Virtual, Augmented and Mixed Reality*, Stephanie Lackey and Jessie Chen (Eds.). Springer International Publishing, Cham, 215–226.

[12] John Brooke. 1995. SUS: A quick and dirty usability scale. *Usability Eval. Ind.* 189 (11 1995).

[13] Edwige Chauvergne, Martin Hachet, and Arnaud Prouzeau. 2023. User Onboarding in Virtual Reality: An Investigation of Current Practices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 711, 15 pages. doi:10.1145/3544548.3581211

[14] World Wide Web Consortium. 2018. *Report from W3C Workshop on Permissions and User Consent*. Technical Report. World Wide Web Consortium. https://www.w3.org/Privacy/permissions-ws-2018/report.html Accessed: March 09, 2025.

[15] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. 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* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 64, 16 pages. doi:10.1145/3411764.3445148

[16] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. 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* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 64, 16 pages. doi:10.1145/3411764.3445148

[17] Thomas Franke, Christiane Attig, and Daniel Wessel. 2019. A Personal Resource for Technology Interaction: Development and Validation of the Affinity for Technology Interaction (ATI) Scale. *International Journal of Human–Computer Interaction* 35, 6 (2019), 456–467. arXiv:https://doi.org/10.1080/10447318.2018.1456150 doi:10.1080/10447318.2018.1456150

[18] John K Haas. 2014. *A history of the unity game engine*. Technical Report. Worcester Polytechnic Institute.

[19] Hilda Hadan, Lydia Choong, Leah Zhang-Kennedy, and Lennart E. Nacke. 2024. Deceived by Immersion: A Systematic Analysis of Deceptive Design in Extended Reality. *ACM Comput. Surv.* 56, 10, Article 250 (May 2024), 25 pages. doi:10.1145/3659945

[20] Hilda Hadan, Derrick M. Wang, Lennart E. Nacke, and Leah Zhang-Kennedy. 2024. Privacy in Immersive Extended Reality: Exploring User Perceptions, Concerns, and Coping Strategies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 784, 24 pages. doi:10.1145/3613904.3642104

[21] Ismat Jarin, Yu Duan, Rahmadi Trimananda, Hao Cui, Salma Elmalaki, and Athina Markopoulou. 2024. BehaVR: User Identification Based on VR Sensor Data. arXiv:2308.07304 [cs.HC] https://arxiv.org/abs/2308.07304

[22] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (Mountain View, California, USA) *(SOUPS '09)*. Association for Computing Machinery, New York, NY, USA, Article 4, 12 pages. doi:10.1145/1572532.1572538

[23] Dave S. Kerby. 2014. The Simple Difference Formula: An Approach to Teaching Nonparametric Correlation1. *Comprehensive Psychology* 3 (2014), 11.IT.3.1. arXiv:https://journals.sagepub.com/doi/pdf/10.2466/11.IT.3.1 doi:10.2466/11.IT.3.1

[24] Yeji Kim. 2022. Virtual reality data and its privacy regulatory challenges: A call to move beyond text-based informed consent. *Cal. L. Rev.* 110 (2022), 225.

[25] Fride Haram Klykken. 2022. Implementing continuous consent in qualitative research. *Qualitative Research* 22, 5 (2022), 795–810. arXiv:https://doi.org/10.1177/14687941211014366 doi:10.1177/14687941211014366

[26] Veronika Krauß, Pejman Saeghe, Alexander Boden, Mohamed Khamis, Mark McGill, Jan Gugenheimer, and Michael Nebeling. 2024. What Makes XR Dark? Examining Emerging Dark Patterns in Augmented and Virtual Reality through Expert Co-Design. *ACM Trans. Comput.-Hum. Interact.* 31, 3, Article 32 (Aug. 2024), 39 pages. doi:10.1145/3660340

[27] Junsu Lim, Hyeonggeun Yun, Auejin Ham, and Sunjun Kim. 2022. Mine Yourself!: A Role-playing Privacy Tutorial in Virtual Reality Environment. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 375, 7 pages. doi:10.1145/3491101.3519773

[28] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.

[29] Shady Mansour, Verena Winterhalter, Florian Alt, and Viktorija Paneva. 2026. A Multi-Layered Privacy Permission Framework for Extended Reality. In *Proceedings of the 2025 New Security Paradigms Workshop (NSPW '25)*. Association for Computing Machinery, New York, NY, USA, 50–65. doi:10.1145/3774761.3774916

[30] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.

[31] Anthony Morton and M. Angela Sasse. 2012. Privacy is a process, not a PET: a theory for effective privacy practice. In *Proceedings of the 2012 New Security Paradigms Workshop* (Bertinoro, Italy) *(NSPW '12)*. Association for Computing Machinery, New York, NY, USA, 87–104. doi:10.1145/2413296.2413305

[32] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2024. SoK: Data privacy in virtual reality. *Proc. Priv. Enhancing Technol.* 2024, 1 (Jan. 2024), 21–40.

[33] Vivek Nair, Gonzalo Munilla Garrido, Dawn Song, and James O'Brien. 2023. Exploring the privacy risks of adversarial VR game design. *Proceedings on Privacy Enhancing Technologies* 2023, 4 (2023), 238–256. doi:10.56553/popets-2023-0108

[34] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 895–910. https://www.usenix.org/conference/usenixsecurity23/presentation/nair-identification

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

[36] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147. arXiv:https://doi.org/10.1080/1369118X.2018.1486870 doi:10.1080/1369118X.2018.1486870

[37] Martin T Orne. 2017. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. In *Sociological methods*. Routledge, New York, NY, USA, 279–299.

[38] Viktorija Paneva, Marvin Strauss, Verena Winterhalter, Stefan Schneegass, and Florian Alt. 2024. Privacy in the Metaverse:. *IEEE Pervasive Computing* 23, 3 (2024), 73–78. doi:10.1109/MPRV.2024.3432953

[39] Viktorija Paneva, Verena Winterhalter, Naga Sai Surya Vamsy Malladi, Marvin Strauss, Stefan Schneegass, and Florian Alt. 2025. Usable Privacy in Virtual Worlds: Design Implications for Data Collection Awareness and Control Interfaces in Virtual Reality. arXiv:2503.10915 [cs.HC] https://arxiv.org/abs/2503.

10915

[40] Sarah Prange, Pascal Knierim, Gabriel Knoll, Felix Dietz, Alexander De Luca, and Florian Alt. 2024. I do (not) need that Feature! – Understanding Users' Awareness and Control of Privacy Permissions on Android Smartphones. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*. USENIX Association, Philadelphia, PA. https://www.usenix.net/conference/soups2024/presentation/prange

[41] Abhinaya SB, Abhishri Agrawal, Yaxing Yao, Yixin Zou, and Anupam Das. 2025. "What are they gonna do with my data?": Privacy Expectations, Concerns, and Behaviors in Virtual Reality. *Proceedings on Privacy Enhancing Technologies* 2025, 1 (2025), 58–77.

[42] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security* (Ottawa, Canada) *(SOUPS '15)*. USENIX Association, USA, 1–17.

[43] Han Shao, Xiang Li, and Guodi Wang. 2022. Are You Tired? I am: Trying to Understand Privacy Fatigue of Social Media Users. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 378, 7 pages. doi:10.1145/3491101.3519778

[44] Mohammad Tahaei, Ruba Abu-Salma, and Awais Rashid. 2023. Stuck in the Permissions With You: Developer & End-User Perspectives on App Permissions & Their Privacy Ramifications. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 168, 24 pages. doi:10.1145/3544548.3581060

[45] Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. 2014. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 91–100. doi:10.1145/2556288.2557400

[46] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Björn Hartmann. 2019. TutoriVR: A Video-Based Tutorial System for Design Applications in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300514

[47] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Bjoern Hartmann. 2020. TransceiVR: Bridging Asymmetrical Communication Between VR Users and External Collaborators. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 182–195. doi:10.1145/3379337.3415827

[48] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 3789–3806. https://www.usenix.org/conference/usenixsecurity22/presentation/trimananda

[49] Hanqiu Wang, Zihao Zhan, Haoqi Shan, Siqi Dai, Maximilian Panoff, and Shuo Wang. 2024. GAZEploit: Remote Keystroke Inference Attack by Gaze Estimation from Avatar Views in VR/MR Devices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security* (Salt Lake City, UT, USA) *(CCS '24)*. Association for Computing Machinery, New York, NY, USA, 1731–1745. doi:10.1145/3658644.3690285

[50] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android permissions remystified: a field study on contextual integrity. In *Proceedings of the 24th USENIX Conference on Security Symposium* (Washington, D.C.) *(SEC'15)*. USENIX Association, USA, 499–514.

[51] Leah Zhang-Kennedy, Maxwell Keleher, and Michaela Valiquette. 2024. Navigating the Gray: Design Practitioners' Perceptions Toward the Implementation of Privacy Dark Patterns. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 97 (April 2024), 26 pages. doi:10.1145/3637374