Who Explains Privacy Policies to Me? Embodied and Textual LLM-Powered Privacy Assistants in Virtual Reality

Vincent Freiberger
freiberger@cs.uni-leipzig.de
Center for Scalable Data Analytics
and Artificial Intelligence (ScaDS.AI)
Dresden/Leipzig, Leipzig University
Leipzig, Germany

Moritz Dresch
mo.dresch@campus.lmu.de
LMU Munich
Munich, Germany

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

Arthur Fleig[*]
arthur.fleig@uni-leipzig.de
Center for Scalable Data Analytics
and Artificial Intelligence (ScaDS.AI)
Dresden/Leipzig, Leipzig University
Leipzig, Germany

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

Abstract

Virtual Reality (VR) systems collect fine-grained behavioral and biometric data, yet privacy policies are rarely read or understood due to their complex language, length, and poor integration into users' interaction workflows. To lower the barrier to informed consent at the point of choice, we explore a Large Language Model (LLM)-powered privacy assistant embedded into a VR app store to support privacy-aware app selection. The assistant is realized in two interaction modes: a text-based chat interface and an embodied virtual avatar providing spoken explanations. We report on an exploratory within-subjects study ($N = 21$) in which participants browsed VR productivity applications under unassisted and assisted conditions. Our findings suggest that both interaction modes support more deliberate engagement with privacy information and decision-making, with privacy scores primarily functioning as a veto mechanism rather than a primary selection driver. The impact of embodied interaction varied between participants, while textual interaction supported reflective review.

CCS Concepts

• **Security and privacy → Human and societal aspects of security and privacy**; • **Human-centered computing → Human computer interaction (HCI)**.

Keywords

Usable Privacy, Virtual Reality, Privacy Policy, LLMs

ACM Reference Format:
Vincent Freiberger, Moritz Dresch, Florian Alt, Arthur Fleig, and Viktorija Paneva. 2026. Who Explains Privacy Policies to Me? Embodied and Textual LLM-Powered Privacy Assistants in Virtual Reality. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3772363.3798567

1 Introduction

VR systems rely on a range of high-fidelity sensing technologies, such as cameras, eye-tracking, and inertial measurement units, to capture user behavior with millimeter-level precision [11, 18, 26]. While these capabilities enable immersive experiences, they also transform subtle movements and gaze patterns into highly identifiable biometric data. As little as 100 seconds of motion data can identify individuals within a pool of 55,541 with 94.33% accuracy [24]. VR data can reveal potentially sensitive inferences such as cognitive load, emotional states, and even sexual orientation [25], creating a form of biometric psychography [17] more revealing than conventional web browsing. Despite these risks, VR privacy documentation remains static and opaque, often buried in external 2D menus that users rarely consult before installation [31, 40]. This lack of transparency creates a critical awareness gap: users significantly underestimate the tracking permissions they grant [37] and the granularity of inferred information, such as emotional or cognitive states [12], increasing their vulnerability to hyper-targeted influence, such as attention-driven targeted advertising [1].

Large Language Models offer a promising approach to addressing this gap. Conversational agents can summarize complex data practices, adapt explanations to user literacy levels, and answer situational questions [9, 35]. These capabilities could address key design challenges for privacy-related user interfaces in VR, such as balancing user engagement with privacy awareness and breaking down privacy policy information for user comprehension [29, 30].

To explore this potential, we embed an LLM-powered privacy assistant directly at the point of decision – a VR app store. The assistant is presented in two interaction modes: a text-based floating chat panel and an embodied virtual avatar providing spoken explanations. Our work is guided by the following research questions:

RQ1 How can LLM-powered privacy assistants impact privacy awareness and decision-making in VR?

[*]Equal last author contribution.

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/2026/04
https://doi.org/10.1145/3772363.3798567

RQ2 How does interface modality (text chat vs. avatar) affect user comprehension, interaction, and overall experience?

We address these questions through an exploratory within-subjects study (N=21), comparing unassisted browsing with the two assisted conditions: a text-based chat panel and an embodied avatar. In semi-structured interviews, we qualitatively investigate participants' decision-making, modality preferences, and tool interaction.

Our findings show that, both the avatar and chat conditions supported participants' perceived privacy risk comprehension and awareness during app selection, with participants more frequently ruling out applications they perceived as higher risk. The avatar supported a more natural and socially engaging interaction for some participants, while others experienced it as distracting or uncanny and preferred the text-based chat interface. The chat modality enabled careful reading and reflection, though some participants reported feeling overwhelmed by the volume of text. Overall, our early results suggest that both embodied virtual agents and text-based chatbots present viable approaches for supporting informed consent in VR, while warranting larger-scale validation across diverse VR contexts and user populations.

Our **contributions** are twofold: 1) **Design and implementation of an LLM-powered privacy assistant integrated into a VR app store**, demonstrating how privacy policy assistance can be embedded directly at the point of decision through both text-based and embodied interaction modalities. 2) **Empirical evaluation of chat-based and embodied privacy assistance**, providing qualitative insights into impact on privacy awareness and user experience, as well as the modality-specific trade-offs.

2 Related Work

Prior work spans empirical risk assessments, transparency gaps, and automated privacy policy assistants.

Biometric Risks. VR headsets expose users to privacy risks distinct from web and mobile ecosystems. Even though generalizability across applications is contested by some research [33, 36], continuous telemetry streams create unique biometric fingerprints: motion trajectories from head and hand sensors can identify individuals within seconds with near-perfect accuracy across sessions [23–25]. Similarly, eye-tracking data has been shown to effectively identify individuals [7] and reveal cognitive load and emotional arousal [26]. Biometric data can reveal sensitive health indicators or personality traits [1]. However, users are largely unaware to the extent of these capabilities, often assuming only physical movement is tracked, overlooking deeper emotional or cognitive inferences [12]. Moreover, biometric signatures derived from VR interaction remain stable over months, rendering traditional anonymization ineffective [11]. Despite this heightened exposure, transparency mechanisms in VR remain inadequate. Privacy disclosures are frequently incomplete [11, 40] and often buried in nested menus that are only accessible after installation [30]. As a result, users are frequently uninformed about data practices [1].

AI-Driven Privacy Assistants. VR environments require mechanisms that make abstract data risks visible and understandable to users [6], yet few assistants operate in VR contexts or prioritize user trust [22]. Beyond VR, research has long sought to automate the analysis of privacy policies. Early approaches leveraged classical natural language processing (NLP) and datasets such as OPP-115 [38] to classify policy text. Tools like *Polisis* [13] and *PriBot* [14] pioneered this space by generating icon-based ratings or retrieving policy segments in response to chat queries. Nevertheless, these systems typically returned raw policy excerpts, leaving the burden of legal interpretation on the user.

Subsequent tools attempted to improve contextuality and usability. *PrivacyInjector* [39] overlaid icons directly onto website elements, while *PrivacyCheck* [27] scored policies against predefined criteria (e.g., GDPR compliance). Although such interventions supported more informed decision-making, they remained largely rigid, as users could not query specific concerns or request simplified explanations for complex data flows.

Recent work has increasingly turned to LLMs to bridge this comprehension gap. Studies show that LLMs match or exceed traditional NLP approaches in extracting data types [32, 41] and can reduce cognitive load by summarizing risks [5, 35]. For instance, Sun et al. demonstrated that an LLM-based agent can effectively categorize and summarize policy sections, increasing user confidence [35]. Similarly, the PRISMe browser extension showed that interactive, LLM-based exploration privacy policies can improve policy understanding and risk awareness during browsing [9].

Summary and Research Gap. While VR systems enable granular and persistent data collection, current transparency mechanisms are insufficient and poorly integrated. While LLM-based privacy assistants show promise for improving policy comprehension, they lack integration into immersive, point-of-decision contexts. It remains unclear how such assistants should be integrated into VR environments or how interaction modalities affect user comprehension, usability, and privacy-related decision-making. To address these gaps, we embed an LLM-powered privacy assistant into a VR app store and compare text-based versus embodied modalities for privacy-aware VR app selection.

3 System Design

We developed a VR application in Unity that simulates an app marketplace, running on a Meta Quest 3 headset. The user is positioned at a virtual desk with a floating, repositionable UI panel in front of them. Selecting an app opens a product card that mirrors standard Quest 3 storefronts, displaying images, description, and pricing information. To minimise confounding effects arising from app genre, each experimental condition included different, existing applications sampled from the same productivity-focused categories (e.g., virtual meetings, virtual desktop, or 3D sketching; see Appendix A). The privacy assistant employs a layered design approach inspired by PRISMe [9], guiding users from a privacy dashboard to either a chat or an avatar to query about specifics, depending on the experimental condition. Here, both interfaces display a set of three context-aware suggestions to support users with limited privacy expertise. The dashboard displays ten "traffic-light" (red-yellow-green) indicators representing specific data categories (e.g., Sensor Access, User Rights). This allows users to identify specific risk factors at a glance.

Users are presented with the following interface variations:

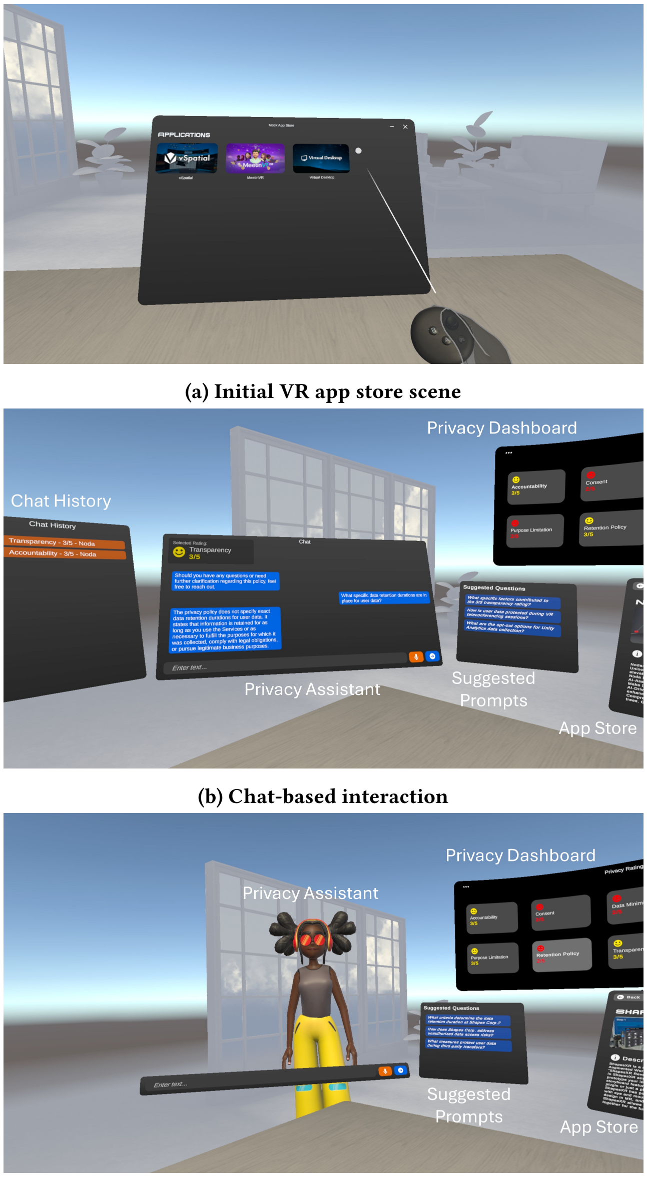

Should you have any questions or need further clarification regarding this policy, feel free to reach out.

What specific data retention durations are in place for user data?

The privacy policy does not specify exact data retention durations for user data. It states that information is retained for as long as you use the Services or as necessary to fulfill the purposes for which it was collected, comply with legal obligations, or pursue legitimate business purposes.

What specific factors contributed to the 3/5 transparency rating?

How is user data protected during VR teleconferencing sessions?

What are the opt-out options for Unity Analytics data collection?

What criteria determine the data retention duration at Shapes Corp.?

How does Shapes Corp. address unauthorized data access risks?

What measures protect user data during third-party transfers?

(a) Initial VR app store scene

(b) Chat-based interaction

(c) Avatar-based interaction

Figure 1: Illustration of the three VR interface conditions used in the user study. (a) ***Baseline:*** **storefront-only app store view (b)** ***Chat condition:*** **text-based interface and privacy ratings. (c)** ***Avatar condition:*** **spoken explanations, gesture cues and privacy ratings.**

Baseline (Storefront Only): The control condition displays only the standard product card (Figure 1a) and descriptions, without the privacy assistant.

Chat Assistant: In this condition, a privacy dashboard appears left of the main store window (Figure 1b upper right). Users can select a rating category to receive a text summary or type/speak free-form questions in a chat panel (Figure 1b middle). A scrollable transcript (Figure 1b left) maintains conversation history for later review.

Avatar Assistant: In this condition, instead of a chat window, a 3D humanoid character (sourced from Mixamo[1]) appears behind the desk (Figure 1c). The avatar utilizes uLipSync [16] for real-time lip-synchronization and procedural gestures. It responds to the same inputs as the chat interface but delivers information via synthesized speech and non-verbal cues.

Both chat and avatar conditions share the same LLM backend and visual design, differing only in their presentation modality.

Backend. For each app, the corresponding privacy policy text was cached locally and provided to the LLM at query time. We utilize OpenAI's GPT-4o API [28] for dashboard data generation and question answering functionality. We borrowed the prompting for generating ratings from PRISMe [10] and adjusted chat prompting to our VR context (see Appendix B). Voice interaction is handled via Meta's Wit.ai service [21] for speech-to-text and text-to-speech conversion.

4 User Study

We conducted a within-subjects lab study to investigate how different privacy interface conditions influence users' privacy awareness understanding, and overall experience. In the *Baseline* condition, participants viewed static app descriptions and standard storefront information without interactive privacy support. The *Chat* and *Avatar* conditions each introduced an integrated LLM-powered privacy policy assistant.

Participants. We recruited 21 participants (9 female, 12 male). The sample included 14 students and 7 employed professionals. Sixteen participants had prior VR experience (primarily gaming), while five were novices. The study lasted approximately 60 minutes, and participants received either €12 or course credits as compensation.

Procedure. After providing informed consent, participants completed a demographic questionnaire, the Affinity for Technology Interaction (ATI) scale [8], and the Internet Users' Information Privacy Concerns (IUIPC) scale [20] to establish technical affinity and baseline privacy attitudes.

Participants then completed three VR sessions. The first session was always the *Baseline* to capture normative unprimed app-browsing behavior, followed by the *Chat* and the *Avatar* conditions in a counterbalanced order. In each session, participants were presented with three productivity apps (different apps in each condition). They were instructed to browse the app store, think aloud, and select one app they would pick to "install on your own device" or reject all options. In both assistant conditions, selecting a privacy rating triggered an auto-generated overview. Participants could then ask unstructured questions via voice input, text input, or suggested prompts until they felt ready to make a decision.

After completing all three VR sessions, participants took part in a semi-structured interviews focusing on their interaction experience, decision making process, and modality preferences (see Appendix C for the interview guide).

Data Analysis. Interview transcripts were analyzed using inductive thematic analysis [3]. Two researchers independently coded the data using Atlas.ti [2] and then met to discuss discrepancies and iteratively refine the codebook, resulting in 16 codegroups (see Supplementary Material).

[1] www.mixamo.com

5 Results

To contextualize our qualitative findings, we report participants' technical affinity and baseline privacy attitudes. Scores on the ATI scale indicated high technical affinity in general ($M = 4.20, SD = 0.80$). Participants also reported relatively high privacy concern across the IUIPC dimensions of *Awareness* ($M = 5.95, SD = 1.31$), *Control* ($M = 5.67, SD = 1.28$), and *Collection* ($M = 5.38, SD = 1.89$).

All participants inspected for each app at least two rating categories and posed an average of 9 follow-up questions. Three categories accounted for 53% of all look-ups: *User Rights* (22%), *Consent* (17%), *Sensor Access* (14%). Next, we present our qualitative results, structured around the four identified main themes.

Balancing Engagement and Control. Some participants found the avatar's presence engaging and natural (P10, P12, P15), while others experienced the anthropomorphism as uncomfortable – for example, P6 *"hated having someone stare at [them]"* – or disliked adapting to its fixed speech pace (P8, P13-14). Participants also reported several usability frictions, including the system's inability to handle natural speech pauses, causing interruptions (P7), having to click on UI elements disrupting conversation flow (P1), and a lack of option to read responses alongside the audio (P2, P6, P7, P21) or review conversation history (P1, P5, P7, P9). The chatbot addressed some of these limitations by supporting review and reflection (P1-2, P5, P7-9), allowing users to *"thoroughly understand and check the content of the answer provided"* (P2). Overall, participants praised the tool's robustness to spelling or pronunciation errors and reliable intent recognition (P1). The tool's intuitive design (P1, P7, P9, P16, P18), being *"easy to understand and [...] feel[ing] like it is for anybody"* (P11) helped users. Overall, balancing engagement and control emerges as a key design challenge, with embodied agents drawing users in while textual interfaces better support deliberate reflection.

Layered Privacy Information and Personalization Needs. Participants generally appreciated the layered design with the dashboard providing a *"one view shot"* overview (P1), and deeper inquiry options facilitated through follow-up questions, which helped users' understanding (P1-2, P10, P21). However, preferences varied: some would have preferred directly posing questions without first navigating the dashboard (P1), while others asked for more detailed justifications for individual ratings upfront (P5, P16, P19-20), including policy excerpts or concrete examples (P8, P16). Perceptions of information density also differed. While some found the amount of information to be *"at a good level"* (P4) or valued them for providing meaningful choice (P10), others experienced the ten dashboard criteria as overwhelming (P5-6). These divergent reactions highlight the need for personalization, including adjustable information granularity and presentation. Additional suggestions included an at-a-glance overall privacy score for easier app comparison (P2, P4, P6, P21), customizable panel layouts (P5, P9, P19), adjustable UI sizing (P14, P19), and configurable VR scenes/backgrounds (P7, P12, P21). Participants emphasized that such a tool would only have meaningful impact if seamlessly integrated into the app store experience like in our prototype (P7, P20).

Raising Privacy Awareness and Understanding. Prior to using the tool, many participants reported indifference towards privacy policy information and described routinely accepting tracking without scrutiny (P1, P7, P9, P14). Interaction with the privacy assistant prompted several users to shift from indifference to a more active evaluation of data practices. For example, P14 realized, *"what they [the app] were tracking and some things I might not have been comfortable with"*. The dashboard's quick overview, combined with question-answering functionality that *"quite consistently added information"* (P2) supported *"informed decision[s]"* (P2), and encouraged participants to consider privacy more deliberately in their decision-making. P1 reflected that, *"we all need to be more cautious, and this panel really helps"*. Several participants further reported intentions to engage more critically with privacy policies in the future, even in the absence of a dedicated privacy policy assessment tool (P1-2, P14). At the same time, one participant reflected on issues of trust in AI-generated explanations, noting the importance of maintaining a critical stance and indicating they would *"read the actual privacy policies and compare them against the chatbot results"* (P6).

Privacy as a Veto Mechanism in App Selection. Participants rarely described privacy as the primary driver for selecting a VR application (P6, P10, P20). Instead, it functioned as a veto mechanism (P6, P10-11, P19). Across conditions, participants prioritized pragmatic and hedonic factors, such as price (P1, P5, P8, P16), app-specific utility (P7-10, P14, P20), and visual appeal (P9, P12-13, P15). Participants described the privacy assistant as particularly useful for assessing apps that lacked *"instinctive trust"* (P10), highlighting its role in trust calibration. Some participants used the question-answering functionality to probe specific concerns, such as the availability of opt-out options (P6). Others reported eliminating options by ruling out apps that received poor privacy ratings on the dashboard (P7, P11, P13, P20).

6 Discussion

We discuss how LLM-powered privacy assistants can impact privacy awareness and decision-making in VR (RQ1) and how interface modality (text chat vs. avatar) affects user comprehension, interaction, and overall user experience (RQ2).

Addressing RQ1: Our findings suggest that the LLM-powered privacy assistant functions primarily as a *supporting decision scaffold* rather than as a primary decision factor. While visual appeal, price, and perceived functionality remained dominant factors, the assistant prompted users to engage more deliberately with privacy-related information and reflect on the associated risks. By embedding privacy support directly into the VR app store and surfacing risks at the point of decision, the assistant reflects core Privacy by Design principles of proactive disclosure and privacy embedded into system design [4]. While users valued the assistant's granular breakdown of privacy criteria, the presence of multiple categories occasionally led to information overload, highlighting the potential value of progressive disclosure and personalization mechanisms that adapt the level of detail to users' needs and momentary attention. Finally, several participants emphasized that the assistant's impact would be contingent on seamless integration into existing VR app store workflows, highlighting that the effectiveness of privacy assistance tools is shaped not only by their information quality, but also by frictionless embedding within users' established interaction routines.

Addressing RQ2: With respect to interface modality, our findings point to a trade-off between engagement and comprehension. The embodied avatar often served as an effective entry point for initiating interest in privacy-related information, with some participants reporting higher engagement and trust, in line with prior work showing that anthropomorphic agents can foster perceived empathy and trust [19]. At the same time, other participants felt uncomfortable in the avatar's presence or constrained by its fixed pacing. In contrast, many participants preferred the text-based chat for understanding complex data practices, emphasizing the ability to re-read and process information at their own pace. This suggests that spoken summaries may be well-suited for initial low-friction exploration, whereas deeper engagement benefits from reviewable, text-based information. Future work could explore hybrid interfaces that combine the avatar's engaging qualities with a textual transcript, allowing users to switch modalities depending on task complexity and preference.

Limitations. Methodically, the lab setting and think-aloud protocol might have heightened participants' privacy awareness, potentially leading to more privacy-conscious behavior than in everyday use. Furthermore, the study focused on productivity applications, and decision-making may differ in social or gaming environments characterized by stronger immersion and social pressure, which might attenuate the impact of privacy cues. Moreover, our sample of participants skewed towards younger, tech-literate participants. Future work could focus on a longitudinal study, include various application genres, and look into whether and how various populations experience different levels of engagement or friction.

While the Baseline condition was always presented first to capture unprimed browsing behavior prior to exposure to privacy assistance, this fixed ordering may have introduced learning or sensitization effects. Future work should therefore randomize or temporally separate baseline exposure to better control for potential carryover effects. Finally, future studies could incorporate objective comprehension and awareness measures to complement self-reported insights this study focused on.

As with AI-driven assistance more broadly, reliance on LLM-based explanations warrants caution, including the risk of hallucination. While this work focused on interaction modality rather than validating the LLM's output, mitigations involving Retrieval Augmented Generation (RAG) and Hallucination Aware Tuning have shown to be effective in reducing hallucinations [10, 34] and could be implemented in future iterations. Prior work on conversational XAI has shown that, while dialogue-based systems can improve understanding, they might also increase over-reliance on AI recommendations [15]. This highlights the importance of designing privacy assistants that support critical engagement and trust calibration.

Our system relied on a remote LLM and a voice interaction service, meaning that user interactions may themselves be subject to external data processing. While this reflects common deployment practices, it introduces a trade-off between explanation quality and the privacy guarantees of the assistant itself. Future work should explore this tension between user privacy (favoring smaller, locally deployed LLMs) and explanation quality (favoring larger, typically commercial models like GPT-4o, which we used).

7 Conclusion

We introduced an LLM-powered privacy assistant embedded in a VR app store and compared unassisted browsing with two assisted conditions, a chatbot and an embodied avatar, in a within-subjects study (N=21). Our findings showed that both forms of assistance encouraged more deliberate engagement with privacy information, with privacy information primarily functioning as a veto mechanism to rule out applications with unacceptable data practices. The embodied avatar fostered interest and social presence for some users, whereas the text-based chat interface better supported careful review and self-paced reflection. This work demonstrates the potential of AI-driven privacy assistants to make privacy considerations more accessible and actionable at the point of decision in virtual immersive environments.

Acknowledgments

This work has received funding from the German Research Foundation (DFG) under grant agreement no. 521584224. Vincent Freiberger and Arthur Fleig acknowledge the financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research „Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI.

References

[1] Melvin Abraham, Pejman Saeghe, Mark Mcgill, and Mohamed Khamis. 2022. Implications of XR on Privacy, Security and Behaviour: Insights from Experts. In *Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) *(NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 30, 12 pages. doi:10.1145/3546155.3546691

[2] ATLAS.ti Scientific Software Development GmbH. 2026. ATLAS.ti: The Qualitative Data Analysis & Research Software. https://atlasti.com

[3] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI research: Going behind the scenes.* Morgan & Claypool Publishers, San Rafael, CA, USA. 51–60 pages. doi:10.2200/S00706ED1V01Y201602HCI034

[4] Ann Cavoukian et al. 2009. Privacy by design: The 7 foundational principles. *Information and privacy commissioner of Ontario, Canada* 5, 2009 (2009), 12.

[5] Chaoran Chen, Daodao Zhou, Yanfang Ye, Toby Jia-Jun Li, and Yaxing Yao. 2025. CLEAR: Towards Contextual LLM-Empowered Privacy Policy Analysis and Risk Generation for Large Language Model Applications. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25).* Association for Computing Machinery, New York, NY, USA, 277–297. doi:10.1145/3708359.3712156

[6] Brendan David-John, Diane Hosfelt, Kevin Butler, and Eakta Jain. 2021. Let's SOUP up XR: Collected thoughts from an IEEE VR workshop on privacy in mixed reality. In *VR4Sec: Security for VR and VR for Security, SOUPS 2021 Workshop.* USENIX Association, USA, 1–6.

[7] Brendan David-John, Diane Hosfelt, Kevin Butler, and Eakta Jain. 2021. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization and Computer Graphics* 27, 5 (2021), 2555–2565. doi:10.1109/TVCG.2021.3067787

[8] 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.

[9] Vincent Freiberger, Arthur Fleig, and Erik Buchmann. 2025. "You Don't Need a University Degree to Comprehend Data Protection This Way": LLM-Powered Interactive Privacy Policy Assessment. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25).* Association for Computing Machinery, New York, NY, USA, Article 36, 12 pages. doi:10.1145/3706599.3719816

[10] Vincent Freiberger, Arthur Fleig, and Erik Buchmann. 2026. Helping Johnny Make Sense of Privacy Policies with LLMs. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (Barcelona, Spain) *(CHI '26).* Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3772318.3791465

[11] Alberto Giaretta. 2024. Security and privacy in virtual reality: a literature survey. *Virtual Reality* 29, 1 (Jan 2024), 1–32. doi:10.1007/s10055-024-01079-9

[12] 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

[13] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G Shin, and Karl Aberer. 2018. Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, USA, 531–548.

[14] Hamza Harkous, Kassem Fawaz, Kang G Shin, and Karl Aberer. 2016. PriBots: Conversational privacy with chatbots. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, USA, 1–6.

[15] Gaole He, Nilay Aishwarya, and Ujwal Gadiraju. 2025. Is Conversational XAI All You Need? Human-AI Decision Making With a Conversational XAI Assistant. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 907–924. doi:10.1145/3708359.3712133

[16] hecomi. 2024. uLipSync: MFCC-based LipSync plug-in for Unity. https://github.com/hecomi/uLipSync. Accessed: 2024-05-22.

[17] Brittan Heller. 2020. Watching androids dream of electric sheep: immersive technology, biometric psychography, and the law. *Vanderbilt Journal of Entertainment & Technology Law* 23 (2020), 1–52.

[18] Julian Kulozik and Nathanaël Jarrassé. 2024. Evaluating the precision of the HTC VIVE Ultimate Tracker with robotic and human movements under varied environmental conditions. arXiv:arXiv:2409.01947

[19] Ning Ma, Ruslana Khynevych, Yunqiang Hao, and Yahui Wang. 2025. Effect of anthropomorphism and perceived intelligence in chatbot avatars of visual design on user experience: accounting for perceived empathy and trust. *Frontiers in Computer Science* 7 (2025), 1531976.

[20] 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.

[21] Meta. 2024. Wit.ai: Natural Language Processing Platform. https://wit.ai Accessed: 2026-01-08.

[22] Victor Morel, Leonardo Horn Iwaya, and Simone Fischer-Hübner. 2025. AI-Driven Personalized Privacy Assistants: A Systematic Literature Review. *IEEE Access* 13 (2025), 160982–161002. doi:10.1109/ACCESS.2025.3609188

[23] Vivek Nair, Gonzalo M Garrido, Dawn Song, and James F O'Brien. 2023. Exploring the Privacy Risks of Adversarial VR Game Design. *Proceedings on Privacy Enhancing Technologies* 4 (2023), 238–256.

[24] 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 *Proceedings of the 32nd USENIX Conference on Security Symposium* (Anaheim, CA, USA) *(SEC '23)*. USENIX Association, USA, Article 51, 16 pages.

[25] Vivek Nair, Louis Rosenberg, James F. O'Brien, and Dawn Song. 2023. Truth in motion: The unprecedented risks and opportunities of extended reality motion data. *IEEE Security & Privacy* 22, 1 (2023), 24–32. doi:10.1109/MSEC.2023.3330960

[26] Mahsa Nasri, Mehmet Kosa, Leanne Chukoskie, Mohsen Moghaddam, and Casper Harteveld. 2024. Exploring Eye Tracking to Detect Cognitive Load in Complex Virtual Reality Training. In *2024 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, Bellevue, WA, USA, 51–54. doi:10.1109/ISMAR-Adjunct64951.2024.00022

[27] Razieh Nokhbeh Zaeem, Safa Anya, Alex Issa, Jake Nimergood, Isabelle Rogers, Vinay Shah, Ayush Srivastava, and K Suzanne Barber. 2020. PrivacyCheck v2: A Tool that Recaps Privacy Policies for You. In *Proceedings of the 29th ACM international conference on information & knowledge management*. Association for Computing Machinery, New York, NY, USA, 3441–3444.

[28] OpenAI. 2024. GPT-4o. https://openai.com/index/hello-gpt-4o/ Accessed: Jul 2024.

[29] Viktorija Paneva, Marvin Strauss, Verena Winterhalter, Stefan Schneegass, and Florian Alt. 2024. Privacy in the Metaverse. *IEEE Pervasive Computing* 20, 4 (Dec. 2024), 5. doi:10.1109/MPRV.2024.3432953 paneva2024ieeepvc.

[30] 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:arXiv:2503.10915

[31] Ophelia Prillard, Costas Boletsis, and Shukun Tokas. 2024. Ethical Design for Data Privacy and User Privacy Awareness in the Metaverse.. In *ICISSP*. SCITEPRESS – Science and Technology Publications, Rome, Italy, 333–341.

[32] David Rodriguez, Ian Yang, Jose M Del Alamo, and Norman Sadeh. 2024. Large language models: a new approach for privacy policy analysis at scale. *Computing* 106 (2024), 1–25.

[33] Lukas Schach, Christian Rack, Ryan P McMahan, and Marc Erich Latoschik. 2025. Motion-based user identification across xr and metaverse applications by deep classification and similarity learning. *arXiv preprint arXiv:2509.08539* (2025).

[34] Juntong Song, Xingguang Wang, Juno Zhu, Yuanhao Wu, Xuxin Cheng, Randy Zhong, and Cheng Niu. 2024. RAG-HAT: A Hallucination-Aware Tuning Pipeline for LLM in Retrieval-Augmented Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Franck Dernoncourt, Daniel Preoţiuc-Pietro, and Anastasia Shimorina (Eds.). Association for Computational Linguistics, Miami, Florida, US, 1548–1558. doi:10.18653/v1/2024.emnlp-industry.113

[35] Bolun Sun, Yifan Zhou, and Haiyun Jiang. 2025. Empowering Users in Digital Privacy Management through Interactive LLM-Based Agents. In *The Thirteenth International Conference on Learning Representations*. International Conference on Learning Representations, Appleton, WI, USA, 1–21.

[36] Qidi J Wang, Alec G Moore, Nayan N Chawla, and Ryan P McMahan. 2024. Cross-Domain Gender Identification Using VR Tracking Data. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 180–189.

[37] Chris Warin, Viktoriya Pak, and Delphine Reinhardt. 2025. Privacy Perceptions Across the XR Spectrum: An Extended Reality Cross-Platform Comparative Analysis of A Virtual House Tour. *Proceedings on Privacy Enhancing Technologies* 1 (2025), 150–168.

[38] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N. Cameron Russell, Thomas B. Norton, Eduard Hovy, Joel Reidenberg, and Norman Sadeh. 2016. The Creation and Analysis of a Website Privacy Policy Corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Katrin Erk and Noah A. Smith (Eds.). Association for Computational Linguistics, Berlin, Germany, 1330–1340. doi:10.18653/v1/P16-1126

[39] Maximiliane Windl, Niels Henze, Albrecht Schmidt, and Sebastian S Feger. 2022. Automating contextual privacy policies: Design and evaluation of a production tool for digital consumer privacy awareness. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18.

[40] Yuxia Zhan, Yan Meng, Lu Zhou, Yichang Xiong, Xiaokuan Zhang, Lichuan Ma, Guoxing Chen, Qingqi Pei, and Haojin Zhu. 2024. VPVet: Vetting Privacy Policies of Virtual Reality Apps. 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, 1746–1760. doi:10.1145/3658644.3690321

[41] Shuning Zhang, Xin Yi, Shixuan Li, Haobin Xing, and Hewu Li. 2025. PrivCAPTCHA: Interactive CAPTCHA to Facilitate Effective Comprehension of APP Privacy Policy. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 385, 20 pages. doi:10.1145/3706598.3713928

A Applications Used in the Study

The following VR productivity applications were used across the experimental conditions.

- Immersed (https://immersed.com)
- Meta Horizon Workrooms (https://forwork.meta.com/de/horizon-workrooms)
- MeetinVR (https://www.meetinvr.com)
- Virtual Desktop (https://www.vrdesktop.net)
- vSpatial (https://www.vspatial.com)
- Noda (https://noda.io)
- Gravity Sketch (https://gravitysketch.com)
- ShapesXR (https://www.shapesxr.com)
- Figmin XR (https://overlaymr.com)

B Prompting

B.1 Dashboard Data

System Prompt: Your output must be a maximum of 600 words long! You are an expert in data protection and a member of an ethics council. You are given a privacy policy. Your task is to uncover aspects in data protection declarations that are ethically questionable from your perspective. Proceed step by step:

(1) Criteria: From your perspective, identify relevant ethical test criteria for this privacy policy as criteria for a later evaluation.

When naming the test criteria, stick to standardized terms and concepts that are common in the field of ethics. Keep it short!

(2) Analysis: Based on this, check for ethical problems or ethically questionable circumstances in the privacy policy.
(3) Evaluation: Only after you have completed step 2: Rate the privacy policy based on your analysis regarding each of your criteria on a 5-point Likert scale. Explain what this rating means. Explain what the ideal case with 5 points and the worst case with one point would look like. The output in this step should look like this:
[Insert rating criterion here]: [insert rating here]/5 [insert line break] [insert justification here]
(4) Contextualization: Use the **attached images and the description of the VR application** to establish a contextual link between the privacy policy and the actual use case of the VR game or experience. Assess whether the data protection measures described in the policy appropriately address the specific risks, user interactions, and immersive data environments involved in the application. Take into account unique concerns such as biometric tracking, gaze monitoring, spatial mapping, and behavioral profiling that are often present in immersive technologies.
(5) Conclusion: Reflect on your evaluation and check whether it is complete.

Important: Check for errors in your analysis and correct them if necessary before the evaluation. You must present your approach clearly and concisely and follow the steps mentioned. Your output must not exceed 600 words.
User Prompt: <Privacy Policy inserted here>

B.2 Prompting Approach for Chat Interface

System Prompt: Privacy policy: <Privacy Policy inserted here> | Rating: <Rating category inserted here> with rating <Rating inserted here>
Users want to know more about how this rating is justified in the privacy policy. When answering the questions, focus on the given topic of the rating. Keep it short!
User Prompt: <User Question inserted here>

B.3 Prompting Approach for Avatar

System Prompt: Privacy policy: <Privacy Policy inserted here> | Rating: <Rating category inserted here> with rating <Rating inserted here>
Users want to know more about how this rating is justified in the privacy policy. When answering the questions, focus on the given topic of the rating. Make sure that the answer is MAX. 25 WORDS! Keep the answer in a conversational style. The output will be transmitted to a TTS interface! Make sure to make full sentences which are understandable!
User Prompt: <User Question inserted here>

B.4 Prompting Approach for our Suggested Question Generation

System Prompt: Your task is to generate thoughtful and critical questions regarding a privacy policy document. These questions should aim to clarify the scope, data handling practices, user rights, or compliance aspects of the policy. Your output must consist of exactly three well-structured questions. Format your output as valid JSON, using the following structure: "suggestions": ["Question 1", "Question 2", "Question 3"] Please ensure the questions are precise, relevant, and helpful for evaluating the transparency and completeness of a privacy policy. Make sure the questions are MAX. 9-14 WORDS!
User Prompt: Specifically: Generate questions about the privacy policy of the app <App Title>, which has a rating of <Rating category inserted here> with a score of <Rating inserted here>. Focus particularly on aspects that led to this rating. Consider critically examining any concerning practices or noteworthy protections mentioned in the following policy: <Privacy Policy inserted here>

C Interview Guide

Overall Experience

- Describe in your own words your experience with the two privacy policy interfaces (Chat and Avatar).

Interface Comparison

- What are some positive and negative aspects of each interface in your own opinion?
- How would you compare the interaction experience with each?
- How well were you able to obtain the information you wanted to know with each one?

Decision Making

- What was your thought process or strategy for choosing the VR apps? Did the approach change between sessions?
- Do you think the privacy assessment tool influenced your decisions? If yes, how?
- How might this experience impact your future decisions when it comes to choosing apps or reviewing privacy policies?

Improvements and Final Reflections

- How would you improve the privacy policy tool or its interface? Are there any features you wish would have been available when browsing the apps?
- Do you have any final thoughts or feedback about your experience in the study?