

User Comprehension and Comfort with Eye-Tracking and Hand-Tracking Permissions in Augmented Reality

Kaiming Cheng
University of Washington
kaimingc@cs.washington.edu

Mattea Sim
Indiana University
matsim@iu.edu

Tadayoshi Kohno
University of Washington
yoshi@cs.washington.edu

Franziska Roesner
University of Washington
franzi@cs.washington.edu

Abstract—Augmented reality (AR) headsets are now commercially available, including major platforms like Microsoft’s HoloLens 2, Meta’s Quest Pro, and Apple’s Vision Pro. Compared to currently widely deployed smartphone or web platforms, emerging AR headsets introduce new sensors that capture substantial and potentially privacy-invasive data about the users, including eye-tracking and hand-tracking sensors. As millions of users begin to explore AR for the very first time with the release of these headsets, it is crucial to understand the current technical landscape of these new sensing technologies and how end-users perceive and understand their associated privacy and utility implications. In this work, we investigate the current eye-tracking and hand-tracking permission models for three major platforms (HoloLens 2, Quest Pro, and Vision Pro): what is the granularity of eye-tracking and hand-tracking data made available to applications on these platforms, and what information is provided to users asked to grant these permissions (if at all)? We conducted a survey with 280 participants with no prior AR experience on Prolific to investigate (1) people’s comfort with the idea of granting eye- and hand-tracking permissions on these platforms, (2) their perceived and actual comprehension of the privacy and utility implications of granting these permissions, and (3) the self-reported factors that impact their willingness to try eye-tracking and hand-tracking enabled AR technologies in the future. Based on (mis)alignments we identify between comfort, perceived and actual comprehension, and decision factors, we discuss how future AR platforms can better communicate existing privacy protections, improve privacy-preserving designs, or better communicate risks.

I. INTRODUCTION

Augmented reality (AR) technologies have reached the cusp of commercial viability, transforming how we interact with the real world, the digital world, and ourselves.¹ Unlike traditional 2D contexts where users interact with content on flat screens, extensive research from industry and academia aims

¹We use the term “AR” to refer to technologies that place virtual content in a user’s view of a real-world environment. Other works may use other terms to refer to the same or related concepts, including mixed reality (MR) and extended reality (XR).

to reinvent how users naturally and smoothly interact with the virtual 3D world. Eye-tracking [2], [20], [43], [64] and hand-tracking [18], [19], [47], [75] are integral to this evolution, enhancing user immersiveness [39], [43] and bringing different yet more intuitive and natural input modalities.

Existing consumer-facing AR headsets, such as Microsoft’s HoloLens 2, Meta’s Oculus Quest Pro, and Apple’s Vision Pro, are already equipped with advanced sensors to perform eye-tracking and hand-tracking. These sensors enable exciting functionalities, such as navigating and interacting with the virtual space using eye gaze [2], [9] and hand gestures [3], [11], [17], or system performance optimizations [13]. Despite the potential benefit these new features bring, existing research has highlighted privacy concerns associated with both eye-tracking and hand-tracking sensors. For instance, the data captured by these devices could be used for inferring sensitive user attributes [66], [72], [96], predicting interest level [44], [100], and revealing user identity [71], [78], [80].

Depending on the system design, AR systems or applications may access the data from these sensors by asking users for permission, or access may be passively enabled by default. End users may grant or deny permission requests based on their expectations of the utility-privacy tradeoff. If users consent to these sensors without fully understanding the associated risks, they may unintentionally expose themselves to privacy violations and security threats [74], [88]. On the other hand, clear communication of the data collection and privacy techniques can effectively increase users’ willingness to adopt new technologies [99]. While the literature on mobile or web platforms is rich, to our knowledge, there have been no empirical studies on permission-granting in the space of Augmented Reality headsets. Thus, our first foundational research question is focused on comprehensively assessing how permission management works on exemplar examples of modern AR technologies:

- **RQ1: Current Landscape.** What is the current technical landscape for eye-tracking and hand-tracking permissions in AR platforms?

For this work, we focus on three leading examples of AR technologies: the Microsoft HoloLens 2, the Meta Oculus Quest Pro, and the Apple Vision Pro. We base our analysis

on experimentation with real devices and publicly-available information. Informed by our findings to RQ1, we next explore the answers to the following two research questions. At a high level, these research questions ask: how do users feel after being presented with the permission dialogs from the HoloLens 2, the Oculus, and the Vision Pro (e.g., how do they feel about their privacy) (RQ2), and do they understand what it means to grant a permission on these devices (e.g., what are the privacy implications of granting permission) (RQ3)?

More precisely, our next two research questions are:

- **RQ2: User Perceptions.** How do people perceive different platforms’ privacy permission flows for eye-tracking and hand-tracking in AR? We explore the extent to which people feel comfortable and informed about these permissions.
- **RQ3: User Comprehension.** After viewing the information provided by the permission flow, how well do people comprehend the permissions, their capabilities, and the associated privacy risk?

To answer RQ2 and RQ3, we conducted a survey of 280 participants. In this survey, we showed participants screenshots of the permission-granting interfaces for the HoloLens 2, Oculus, and Vision Pro. We asked participants to what extent they felt comfortable and informed about the permission, confident about the protection of sensitive data, and how clear they found the permission flow to be. We explicitly recruited participants who had no prior experience with AR, in order to capture people’s comfort and comprehension on their *first* exposure to these permission-granting flows, rather than relying also on their past experiences.

Among our findings, we observe that: (1) the extent to which participants felt comfortable and informed depended on the device, sensor, and whether they were considering system-level or app-level access (Section V-A). (2) Participants experienced greater difficulty understanding privacy implications compared to utility, and are generally less informed at the app-level compared to the system-level (Section V-B). (3) Participants were largely uninformed about data handling processes, for example, whether the system or application shares their data with external servers, has access to the raw data, or accesses their data in the background (Section V-C).

Additionally, we investigate what factors participants report would contribute to their willingness to try eye- and hand-tracking enabled AR technologies (RQ4, Section V-D). For example, how do participants weigh the importance of understanding who has access to their data or why these data are being collected?

- **RQ4: Factors that Impact User Decisions.** What permission-related factors do people report as important in their decision-making process around whether or not to try eye- and hand-tracking enabled AR technologies in the future?

Stepping back, we then compare the results between perception (RQ2), comprehension (RQ3), and self-reported decision factors (RQ4) to identify (mis)alignments (Section VI). For

example, we identify cases where comfort may be founded in part in a misunderstanding of the actual implications or implementation of a permission, meaning that people may believe a permission is more or less privacy-invasive than it actually is. We discuss how future AR systems could improve the permission-granting flow for eye-tracking and hand-tracking while better communicating privacy implications, and/or implementing privacy protections currently lacking.

Disclosure. We have reported all of our findings to Apple, Meta, and Microsoft.

II. BACKGROUND AND RELATED WORK

A. Augmented Reality

AR technologies are receiving increasing attention from both academia and industry. AR headsets like Microsoft’s HoloLens 2 [1], Meta’s Quest Pro [28], and Apple’s Vision Pro headset [8] are transforming previous visions for AR into market-ready products. Unlike traditional mobile computing, AR interactions tailor the immersive digital world in response to a user’s actions, such as eye and hand movement. Today’s consumer AR headsets are equipped with sensors that collect, monitor, and analyze this data in real-time.

B. Eye-Tracking and Hand-Tracking

The current approach to enabling eye- or hand-tracking on AR headsets is through a combination of built-in sensors and computer vision algorithms. For eye-tracking, AR headsets use near-eye infrared cameras [2], [8], [21] to pinpoint the pupil’s location and reveal where the user is looking. For hand-tracking, they use inside-out depth cameras [8], [18], [21] to detect the configuration of each finger as well as hand movement and orientation.

1) *Utility:* Captured eye-tracking and hand-tracking data enable a wide range of AR interactions and offer practical benefits. For example, users can use their hand or gaze as the main input medium to select, navigate, and interact with virtual objects [43], [47], [50], [87]. With recent hardware and computer vision algorithm advancements, an AR system is able to perform high-fidelity 3D reconstruction of a user’s eye or hand movement based on the tracking data, useful for creating social interaction [92], [95]. Eye-tracking data, specifically, enables foveated rendering [13], which optimizes the computational efficiency in rendering by reducing resolution in the periphery, and mitigates the vergence-accommodation conflict [65], which reduces user discomfort. Additionally, today’s AR system can utilize iris patterns for authentication purposes [9], [27].

2) *Privacy Concerns:* Although these sensors have functional benefits as outlined above, recent research has also highlighted the privacy implications associated with eye-tracking and hand-tracking data. For example, prior studies suggest eye-tracking data can be used to reveal sensitive user attributes, including gender, age, race, geographic origin, and a wide array of personal characteristics and preferences [66], [72], [96]. Gaze pattern can be leveraged for targeted marketing

based on a user’s estimated interest level [44], [100]. Recent studies showed users can be profiled and deanonymized based on their hand-tracking data [71], [78], [80]. Because eye-tracking and hand-tracking have such strong implications for privacy, it is integral to investigate the extent to which users understand the capabilities and risks of these sensors.

C. Permission Granting

Users rely on dialogs in the permission-granting process to learn about the potential utility and privacy risks associated with certain permissions, all of which allow users to make an informed decision. Many previous works aim to understand what concerns users have when granting permissions [41], [52]–[55], [77], and how to better design the permission/warning dialog to increase transparency for the users [56], [97]. Prior work assessing the efficacy of permission systems have used comprehension to determine the extent to which users are informed about the permissions being requested. Felt et al. [55] first studied the effectiveness of Android install-time permission, and Shen et al. [86] investigated users’ comprehension of the runtime permission model on iOS and Android. Both studies showed only a very small percentage of users can infer the correct scope of permission capabilities from the system-provided information. Harborth et al. [59] evaluated user comprehension of permissions requested in mobile AR applications. Their results suggested that users are concerned with current permissions in AR, such as speech and face recognition, yet the mobile system did not request permission to collect such data. Our study builds upon prior work by focusing on eye-tracking and hand-tracking permissions, which are unavailable in the mobile AR context and thus are largely novel to much of the population.

D. Security and Privacy Research on AR

We add to a growing body of work from the computer security and privacy community, which has been addressing security, privacy, and safety risks in AR for over a decade [83]. The initial security threat modeling taxonomies for AR were proposed by Roesner et al., identifying input, data access, and output as key areas of concern [84]. Guzman et al. then built on these categories, incorporating user interaction and device protection [48]. Much prior work falls into these taxonomies, including studies focusing on sensor data input privacy in AR/VR platforms [61], [62], [89], [94], [98], device and network safety [58], [93], user input [45], [69], and malicious AR output [38], [46], [67].

Relatedly, other emerging and ongoing research investigates end-user’s privacy preferences within AR. For instance, Denning et al. [49] found that bystanders of AR headset-users are concerned about being identified, highlighting the need to grant permission before being included in the AR recording. O’Hagan et al. [79] conducted an online survey to examine bystanders’ privacy preferences and comfort with various AR functionalities on hypothesized AR applications. Lebeck et al. [68] conducted hands-on HoloLens activities and interviews to examine privacy concerns in multi-user AR environments.

Rauschnabel et al. [81] underscored privacy’s impact on user decision-making in the context of conceptual AR devices. Gallardo et al. [57] explored preferences regarding data collection via hypothetical consumer-grade AR glasses.

However, to the best of our knowledge, no previous research has addressed eye- and hand-tracking permission-granting process of real consumer-level AR devices. Indeed, only recently have such AR headsets and functionalities become available in the consumer market. With the potential for widespread adoption of these headsets in the near future, it is imperative to examine their privacy design, evaluate their permission models, and probe users’ comprehension and potential misconceptions about the technology. We aim to begin closing this gap in our work here.

III. RQ1: CURRENT AR PERMISSION GRANTING LANDSCAPE

A. Methodology

To understand the current landscape of eye-tracking and hand-tracking permissions in today’s AR platforms, we investigated three high-profile publicly available platforms: HoloLens 2 (from Microsoft), Quest Pro (Oculus, from Meta), and Vision Pro (from Apple). Our team conducted multiple rounds of structured brainstorming to generate and refine properties relevant to eye- and hand-tracking permission granting (e.g., whether applications have access to eye-tracking data when running in the background).

After finalizing the properties, the lead author examined the documentation and the privacy policies, and built applications on each device to evaluate each property. We performed our initial evaluation in October 2023 and verified them on the up-to-date AR operating system (Holographics version 24H1, Oculus Quest version 65, and visionOS version 1.1) in May 2024. All paper authors iteratively validated the findings and resolved disagreements.

We highlighted that our findings are based on snapshots of the ever-changing AR permission ecosystem, and the results might be subject to change in future upgrades. For example, we noticed several changes in the permission UI for Oculus hand-tracking and eye-tracking privacy notice, though these changes didn’t affect system capability. Nevertheless, our findings can serve as a benchmark to evaluate how the permission landscape evolves. We summarized the selected properties in the “Permission Comprehension” column in Table IV, highlighting our findings for each AR platform using an underline. Below, we present the key similarities and differences across these three platforms. The complete list of reasoning and supporting references is available in Appendix IX.

B. Eye-Tracking Permission

Permission Request. We find that only Oculus requests the user’s permission to perform eye tracking on a system level, as shown in Figure 1. The permission dialog from the system illustrates the potential utility of eye-tracking and the privacy-preserving techniques Oculus deploys. In contrast,

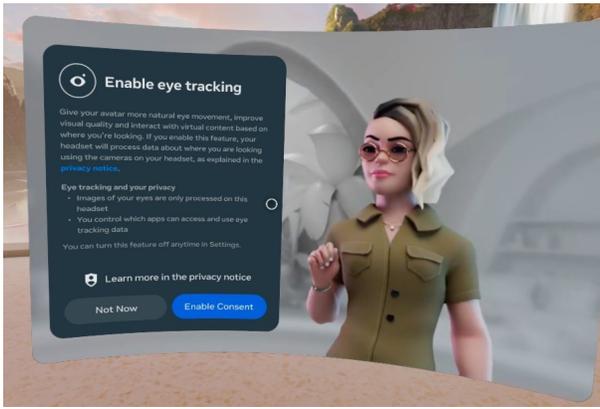
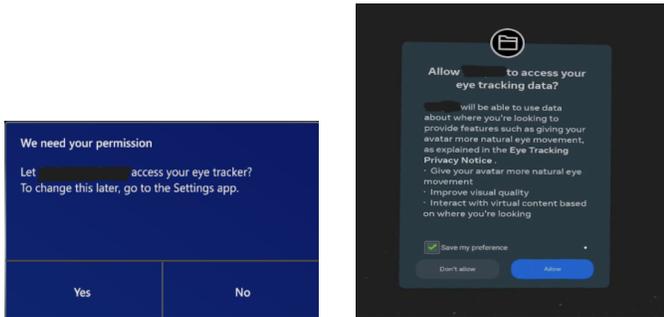


Fig. 1: Oculus: System-level eye-tracking permission.



(a) HoloLens 2: App-level permission (app name blurred for anonymity)

(b) Oculus: App-level permission (app name blurred for anonymity)

Fig. 2: App-level Eye-tracking dialogs

eye-tracking capability is enabled by default for HoloLens or Vision Pro on the system level, given it’s one of the primary input modalities (as opposed to controllers for Oculus). Developers could request eye-tracking permission on Oculus and HoloLens as shown in Figure 2, but not on Vision Pro.

Data Granularity. All three platforms prevent applications from accessing raw eye-tracking images due to significant privacy concerns. For Oculus and HoloLens, the provided eye-tracking APIs [2], [12] include abstracted eye-tracking data, comprising a stream of gaze vectors to represent the user’s eye orientation and movement patterns [14], [20]. However, neither platform controls how third-party entities use, store, or share users’ abstracted gaze data [14].

Compared with Oculus and HoloLens, Vision Pro employs a different, arguably more privacy-preserving, data collection model. According to their Privacy Overview report [31], Apple acknowledges that (abstracted) eye-tracking data, including the content the user looked at or the duration they looked at it, could potentially reveal a user’s thought processes. As a result, while Vision Pro enables eye-tracking permission by default, the processed eye-tracking data is not available to Apple, third-party entities, or websites. Instead, developers utilize Apple’s native event-handling mechanisms, such as UIKit [36] or SwiftUI [35], to manage user interactions automatically. As

users navigate applications, visionOS processes and renders visual effects that respond to where they look on the device.

Data Transmission. While Oculus is the only platform that requests permission to enable eye-tracking on a system level, we also find that it is the only platform to collect and retain user’s eye-tracking data. Specifically, Oculus stored the abstracted gaze data and users’ interactions with eye tracking in their company server. As stated in their privacy policy [14], the eye-tracking data will be associated with users’ accounts until Meta “no longer need it to provide the service or improve the eye-tracking feature”.

C. Hand-Tracking Permission

Permission Request. Similar to eye-tracking, only Oculus requests the user’s permission to perform hand-tracking on a system level, as shown in Figure 3. The permission dialog illustrates the potential utility and provides a reference link to the privacy policy. Vision Pro is the only platform that requests app-level permission for hand-tracking, as shown in Figure 4b, whereas the other two platforms automatically grant applications access to the hand-tracking API. The only platform that supports background access for hand-tracking is HoloLens, as shown in Figure 4a.

Data Granularity. All platforms provide an abstract representation of the user’s hand-tracking through hand skeleton data. With the underlying recognition model, the system can understand users’ gestures, hand position, relative hand size, and hand movement. The only difference is that the developers can get access to the user’s hand-tracking data without an additional prompt on HoloLens and Oculus (if the user already granted it to the system). For Vision Pro, the hand-tracking data is only available to the developer when the application is in an immersive space [19].

Data Transmission. While Oculus is the only platform that requests permission to enable hand-tracking at the system level, it also processes and shares the hand-tracking data with the Oculus server, where it is retained for 90 days [15]. For HoloLens, the hand-tracking data is processed on the device and is not stored [33] and for Vision Pro, the hand-tracking data is only stored on-device [31].

IV. USER STUDY METHODOLOGY

To answer RQ2-RQ4, we designed and ran a user study.

A. Survey Design and Procedure

Inspired by the different permission-granting processes across different sensors and platforms we documented in Section III (RQ1), we designed a survey to study users’ comfort, the extent to which users perceive themselves as informed by the permission granting processes, their comprehension of the permissions, and what factors impact their likelihood of using these devices in the future. This survey, launched online on Prolific in May 2024, assessed perceptions of three AR platforms, with questions designed to answer our research

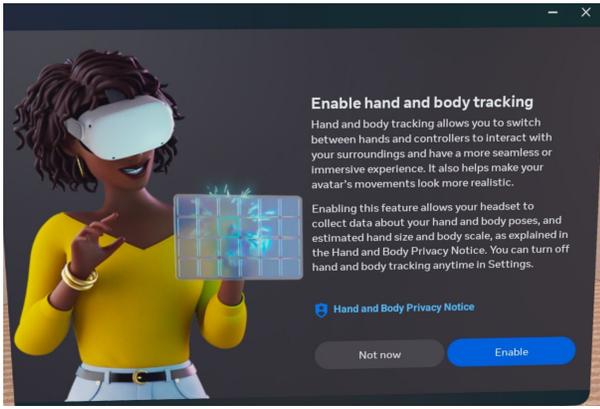
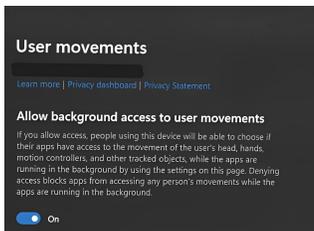
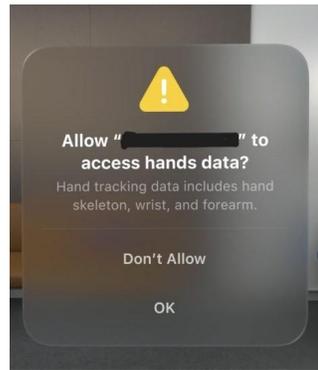


Fig. 3: Oculus: System-level hand-tracking permission.



(a) HoloLens: Background permission for applications. (system name blurred for anonymity)



(b) Vision Pro: App-level permission (app name blurred for anonymity)

Fig. 4: Hand-tracking permission dialogs

questions of interest. The complete list of survey questions and instructions are available in Appendix VIII.

After consenting to participate, participants read that we were investigating perceptions of augmented reality technologies. Participants saw several image examples of AR headsets, and were asked about their familiarity and experience with AR headsets, both broadly and with the headsets investigated in this study specifically. Participants were excluded from analyses (but still received payment) if they indicated they had used any of the three headsets investigated here. Next, participants read that AR headsets have different sensors recording data while the headsets are in use, and that users typically view permission dialogs prompting them to allow or deny the headset access to these data. Participants were told they would view permission dialogs and rate their impressions for two different sensors. Participants were randomly assigned on a between-subjects basis to evaluate one of three mainstream AR headsets: Meta’s Quest Pro, Microsoft’s HoloLens 2, or Apple’s Vision Pro. The company and device names were anonymized in the survey to avoid biasing evaluations. Participants evaluated the device’s eye-tracking and hand-tracking permissions in random order on a within-subjects

basis.

For each sensor, participants were asked to imagine they were using an AR/MR headset with the sensor feature. First, participants read that they were navigating the system-level permission settings for a given sensor. In general, this was followed by a real screenshot of the platform’s permission dialog, or several dialogs depending on the platform’s interface, with all screenshots accompanied by alt text. We also presented other screenshots to simulate the experience of enabling eye- or hand-tracking, such as the hand visualizations that users see when they put on the headset. For platforms that did not explicitly ask for the user’s permission for a given sensor, we told participants that the permission was enabled by default. This part of the survey was designed to follow a user’s actual permission-granting process within a given platform as closely as possible. See Appendix VIII for screenshots.

To assess the extent to which people feel comfortable and informed while experiencing the permission flow (RQ2), participants answered several questions about their perceptions of the dialogs and the device more broadly. Participants responded to a series of 5-point Likert scale questions assessing how informed they felt about both the utility of the permission and its associated privacy risks, their confidence that their data will be securely stored, the extent to which they know what data will be collected and how it will be used based on the permission screenshots presented, and how comfortable they felt using the device (see full questions and scales in Appendix VIII).

We then sought to explore whether the interfaces impacted users’ actual understanding or misperceptions of the system’s capabilities and privacy protections (RQ3). Participants responded to a series of True or False questions about the system’s capabilities and privacy (e.g., “The system can identify which real-world objects you are looking at;” “The system can retain the image of your hand on the AR/MR headset”). For each statement, participants indicated whether they believed it was True or False, or indicated “I don’t know.” Our team conducted multiple rounds of interactive brainstorming and preliminary experiments to generate questions and finalize answers. These questions are inspired by prior studies on mobile permissions (e.g., [55]).

After answering the above questions, participants were then told to imagine they were opening an app on the headset to navigate the app-level permission settings for the device. Here again, participants saw screenshots of permission dialog(s), or received alternative information about the permissions as applicable. Participants responded to the same questions as for the system-level, assessing comfort with the app, how informed they feel, and a similar series of true/false/I don’t know questions about the app’s capability and privacy protections.

Finally, participants read that we wanted to understand what information about the system and app would help them feel more comfortable using this technology in the future. Participants were shown five factors relevant to permission dialogs (i.e., knowing who will have access to the data, how the data will be stored, how the data will be transmitted, what

Gender	Age		Race/Ethnicity		
Man	48.2%	18-24	7.4%	White	87.5%
Women	47.5%	25-34	26.2%	Black or African American	4.6%
Undiscl.	4.3%	35-44	21.4%	Asian	2.5%
		45-54	21.0%	American Indian / Alaskan Native	1.4%
		55-64	14.4%	Native Hawaiian / Pacific Islander	0.4%
		65+	9.6%	Mixed	0.4%
				Undisclosed	3.2%

TABLE I: Breakdown of participant demographics by gender, age, and race/ethnicity.

type of data will be collected, and the purpose of collecting the data). Participants selected their top three most important factors (in no particular order).

Participants answered all questions for a given sensor before evaluating the next sensor. After evaluating both sensors, participants responded to an attention check, reported demographic information, and received payment through Prolific.

B. Ethics

The study was deemed Exempt by the university’s Human Subjects Review Board (IRB). Participants were anonymous and identifying data were removed or not obtained. Participants could leave the survey at any time. Participants were compensated based on Prolific’s guidelines (see below).

C. Participants

We conducted an a priori power analysis using G*Power to determine how many participants were needed to detect a moderate effect size. This analysis determined that 260 participants would be sufficient to detect an effect size of $d = 0.35$ at 80% power in an independent-samples t-test. This sample size also provides sufficient power to detect effect sizes of $n_p^2 < .010$ in mixed-model ANOVAs.² In actuality, 292 adult U.S. crowdworkers on Prolific completed the 13-minute survey online in exchange for payment, with compensation set based on Prolific’s guidelines (\$12 hourly rate). We excluded participants from analyses who failed to pass an attention check and who indicated they had used either Oculus, HoloLens, or Vision Pro. Participants excluded from analyses still received payment. After exclusions, our analyses includes 280 participants. Participants’ demographics are included in Table I.

D. Limitations

We consider several limitations of our study’s design. First, a survey with screenshots may not fully capture the complete experiences of a user wearing an AR headset. Beyond the different modality, there may also be additional information in the device’s initial setup flow, such as a 3D video, that helps communicate permission-related impressions to users that are not captured by our survey design. Similarly, app developers can customize the permission dialog text on Vision Pro or provide justifications before the dialog on Oculus and

²These data do not meet all normality assumptions for ANOVAs. However, prior work shows that ANOVAs are robust against non-normality when the sphericity assumption is met, as it is in our data [42].

HoloLens, meaning that the information shown in the app-level dialog may depend heavily on that customization in practice. Second, our attempt to anonymize company and device names in the survey may not have always been successful. Since certain UI characteristics are manufacturer-specific, they may have been recognizable to some participants, influencing their perceptions. Third, our analysis of participant comprehension depends on our own understanding of the correct answers to the true/false questions (see Appendix IX for our understanding). Nevertheless, we believe it is valuable to understand what participants believe the answers are based on the permission dialogs they see as this understanding will influence user perception and decisions. Lastly, permission designs are subject to change as platforms evolve and update their SDKs. The observations and analyses presented here are based on our understanding of the systems in May 2024. Despite these limitations, our study sheds light on people’s perception and comprehension of novel AR platform permissions and evaluates key aspects of the current designs of these platforms’ permission models and dialogs. Future work must continue to revisit these questions as the technology and app ecosystems evolve, just as a decade or more of research studied the smartphone permission and app ecosystem.

V. RESULTS

We investigated perceptions of eye-tracking and hand-tracking based on the permission flow (RQ2), comprehension of utility and privacy implications (RQ3), and information deemed particularly important to include in the permission dialog (RQ4).

A. RQ2: Perceptions of Permission Flows Differ Across Devices, Sensors, and Use Level

We investigated the extent to which participants felt comfortable and informed using the AR headset. In the sections below, we explore how participants’ perceptions depended on the device and sensor type. Thus, we conduct a series of mixed-methods ANOVAs and t-tests on each dependent variable. We focus on system-level perceptions to avoid inflating Type I errors with additional comparisons at the app-level.

1) *Comfort*: We conducted a mixed-method ANOVA on participants’ comfort level with sensor type (eye-tracking, hand-tracking) as a within-subjects variable and device (Oculus, HoloLens, Vision Pro) as a between-subjects variable. Participants’ comfort was impacted by both the device and the sensor type, indicated by a significant interaction between sensor type and device, $F(2, 277) = 16.108$, $p < .001$, $n_p^2 = .104$ (see all system-level comparisons in Figure 5).

We conducted t-tests across devices and sensors to decompose this interaction. We first observed differences in comfort across devices. In the context of eye-tracking, participants felt similarly comfortable using Oculus and HoloLens, ($p = .387$, $d = 0.13$), but felt significantly more comfortable using both Oculus and HoloLens as compared to Vision Pro ($ps < .009$, $ds > 0.38$). In the context of hand-tracking, participants felt significantly more comfortable using Oculus compared to both

HoloLens ($p < .001$, $d = 0.63$) and Vision Pro ($p = .040$, $d = 0.32$). Participants also felt significantly less comfortable using HoloLens than Vision Pro for hand-tracking ($p = .034$, $d = -0.31$).

Differences in comfort between the sensors, on the other hand, emerged only within HoloLens. Participants who saw dialogs from Oculus or Vision Pro were similarly comfortable with eye-tracking and hand-tracking ($ps > .150$, $ds < .16$). But participants who saw HoloLens dialogs felt significantly more comfortable with the eye-tracking sensor than the hand-tracking sensor ($p < .001$, $d = 0.58$).

2) *Feeling Informed about Permission Utility*: We next investigated the extent to which participants felt informed about the utility of the permissions. At the system level, there was a significant interaction between device and sensor type, $F(2, 277) = 11.394$, $p < .001$, $\eta_p^2 = .076$. For eye-tracking, participants felt similarly informed about the utility of Oculus and HoloLens ($p = .694$, $d = 0.06$). However, participants felt significantly more informed about the utility of both Oculus and HoloLens as compared to Vision Pro ($ps < .040$, $ds > 0.30$). For hand-tracking, participants felt significantly more informed about the utility of Oculus than HoloLens ($p < .001$, $d = 0.63$). There was no significant difference between Oculus and Vision Pro ($p = .122$, $d = 0.24$). Participants felt significantly less informed about the utility of HoloLens compared to Vision Pro ($p = .013$, $d = -0.36$). We next compared differences on the system-level in the extent to which people felt informed about the utility across eye-tracking and hand-tracking. For both Oculus and HoloLens, participants felt significantly more informed about the utility of eye-tracking as compared to hand-tracking ($ps < .030$, $ds > 0.24$). This difference was non-significant amongst participants who saw Vision Pro ($p = .320$, $d = 0.11$).

3) *Feeling Informed about Privacy*: We next investigated the extent to which participants felt informed about the associated privacy risk of the permissions. Once again, at the system-level, there was a significant interaction between device and sensor type, $F(2, 277) = 4.027$, $p = .019$, $\eta_p^2 = .028$. In the context of eye-tracking, participants who saw Oculus felt more informed about the privacy risks than participants who saw Vision Pro ($p = .031$, $d = 0.33$), but no other device comparisons were significant ($ps > .200$, $ds < 0.18$). In the context of hand tracking, participants who saw Oculus felt more informed about the privacy risks than participants who saw either HoloLens or Vision Pro ($ps < .035$, $ds > 0.32$), whereas participants in the latter two conditions did not significantly differ ($ps = .184$, $d = -0.19$).

Comparing across sensors at the system-level, we found that participants who saw both Oculus and HoloLens felt more informed about the privacy risks of eye-tracking than hand-tracking ($ps < .022$, $ds > 0.24$). This difference was non-significant amongst participants who saw Vision Pro ($p = .103$, $d = 0.18$).

4) *Confidence in Security*: We investigated how confident participants felt about the system's ability to securely store their data. There was a significant interaction between device

and sensor type, $F(2, 277) = 4.888$, $p = .008$, $\eta_p^2 = .034$. For eye-tracking, participants felt more confident about Oculus than Vision Pro ($p = .048$, $d = 0.30$), and all other comparisons were non-significant ($ps > .110$, $ds < 0.24$). For hand-tracking, participants felt more confident about Oculus than HoloLens ($p = .011$, $d = 0.37$), and all other comparisons were non-significant ($ps > .190$, $ds < 0.20$). Comparing across sensors at the system-level, participants who saw either Oculus or HoloLens felt more confident in the system securely storing their eye-tracking data than their hand-tracking data, ($ps < .015$, $ds > 0.27$). There was no difference in confidence across sensors for participants who saw Vision Pro ($p = .334$, $d = 0.11$).

5) *Data Use Clarity*: Finally, we investigated the extent to which participants felt they knew what data would be collected and how it would be used (i.e., data clarity) based on the permission flow. At the system-level, there was a significant interaction between device and sensor type, $F(2, 277) = 10.376$, $p < .001$, $\eta_p^2 = .070$. In the context of eye-tracking, there was no significant difference in data clarity across Oculus and HoloLens ($p = .817$, $d = 0.03$). However, participants felt more data clarity from both Oculus and HoloLens as compared to Vision Pro ($ps < .022$, $ds > 0.33$). In the context of hand-tracking, participants felt more data clarity from Oculus as compared to both HoloLens and Vision Pro ($ps < .008$, $ds > 0.40$), and participants in the latter two conditions did not significantly differ ($p = .141$, $d = -0.22$).

Comparing across sensors, across all three devices, participants felt more data clarity about the eye-tracking permission than the hand-tracking permission ($ps < .026$, $ds > 0.24$).

6) *Relationship Between Comfort and Feeling Informed*: The findings above clearly demonstrate that the extent to which participants feel comfortable, informed, and confident are impacted by the permission dialogs and the sensor tracking data in nuanced ways. At a higher level, we were also interested in whether participants who feel more informed also feel more comfortable using the device. Collapsed across all devices and sensors, feeling informed about the utility of the permission ($r = .628$, $p < .001$) and feeling informed about the associated privacy risks of the permission ($r = .595$, $p < .001$) were each correlated with comfort using the device or app. This correlation underscores the importance of *felt* comprehension. Similarly, the extent to which people felt confident that the device was securely storing their data ($r = .792$, $p < .001$) and felt clear about the data use policies ($r = .668$, $p < .001$) were also each correlated with comfort using the device or app. Regardless of actual understanding, feeling more informed and confident after reading permission dialogs may create a more comfortable experience for users — though not necessarily a more privacy-preserving one.

B. RQ3: Permission Comprehension Overview

Users can only make informed security and privacy decisions if they understand the implications of those decisions. The trade-offs between utility and privacy represent the benefits and risks inherent in these choices. Hence, it is

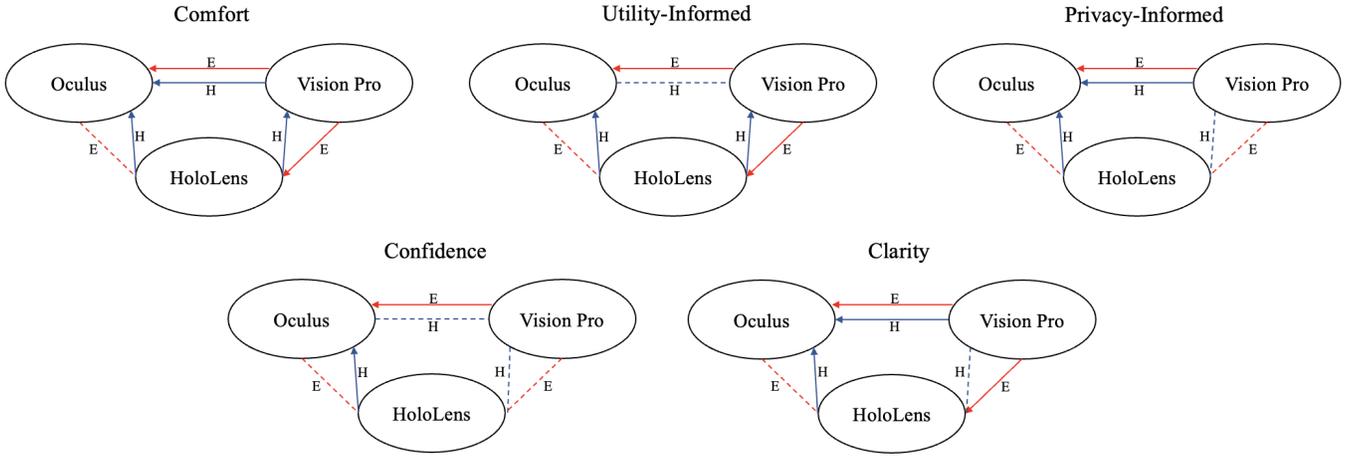


Fig. 5: System-level perceptions compared across devices. Red lines (labeled “E”) represent eye-tracking and blue lines (labeled “H”) represent hand-tracking. Arrows point to the device that was rated significantly higher on the item. Dashed lines = non-significant.

TABLE II: Participant comprehension correctness summary. Hol is HoloLens, Oc is Oculus, Vis is Vision Pro, Avg is the performance on each category, Avg-S is the performance on each sensor, and Avg-T represents the overall performance across all questions.

		Hol	Oc	Vis	Avg	Avg-S	Avg-T
Eye	Util	54.7%	59.8%	43.2%	52.8%	42.6%	43.3%
	Priv	30.2%	45.5%	21.7%	32.4%		
Hand	Util	56.0%	62.1%	56.3%	58.0%	44.0%	
	Priv	30.3%	31.0%	28.5%	30.0%		

TABLE III: Comparing the level of comprehension regarding the system-level permission and application-level permission. See Table IV for details.

	Category	System	App	Diff
Hololens-Eye	Utility	73.0%	36.4%	-36.6%
	Privacy	30.9%	29.5%	-1.4%
Hololens-Hand	Utility	57.1%	54.9%	-2.2%
	Privacy	40.2%	32.8%	-7.4%
Oculus-Eye	Utility	67.9%	51.7%	-16.2%
	Privacy	49.5%	41.4%	-8.1%
Oculus-Hand	Utility	62.9%	61.4%	-1.5%
	Privacy	40.2%	21.7%	-18.5%
Vision-Eye	Utility	70.7%	15.6%	-55.1%
	Privacy	17.6%	25.9%	+8.3%
Vision-Hand	Utility	56.6%	56.1%	-0.5%
	Privacy	26.3%	30.7%	+4.4%

crucial that systems are designed to clearly communicate these factors, enabling users to navigate this balance with clarity and knowledge. In addressing RQ3, we investigate this dynamic by analyzing comprehension differences (1) across various sensors, (2) between system-level and app-level permissions, and (3) among different devices. Appendix IX provides the “answer key”, to the best of our knowledge.

1) *Comprehension Across Sensors*: We first scored participants’ answers to the true/false questions. We found that participants had a slightly better understanding of hand-tracking (average 44.0% across three platforms) than eye-tracking (average 42.6% across three platforms). Although participants generally understood the utility of eye-tracking and hand-tracking, on average scoring 52.8% and 58.0% on utility-related questions respectively, their understanding of privacy implications was noticeably lower. Specifically, participants only correctly answered an average of 32.4% of the privacy questions for eye-tracking and 30.0% for hand-tracking.

2) *Comprehension Across System-Level and App-Level Permissions*: We explored whether respondents’ comprehension differs between system-level permissions and app-level permissions, where we see different technical and UX designs. As shown in Table III, in all conditions examined, participants tended to be less informed regarding the utility of permissions within the application compared to their understanding of the same permission within the system. We observe a sharp decline in the understanding of eye-tracking utility at the app level for HoloLens (a decrease of 36.6%) and Vision Pro (a decrease of 55.1%). For participants’ comprehension of privacy, we observe a similar declining pattern in the understanding of eye-tracking and hand-tracking privacy at the app level for both HoloLens and Oculus. The only exception is Vision Pro, where app-level privacy comprehension is better than system-level.

3) *Comprehension Across Devices*: Lastly, we assess whether users’ comprehension differs across the three devices’ permission-granting flows. Table II summarizes the comprehension score across devices. We conducted two-sample Z-tests to compare across devices. Participants who saw Oculus had a significantly higher comprehension of the eye-tracking utility ($z = 2.1842$, $p = .029$) and privacy ($z = 3.3082$, $p < .001$) compared to participants who saw Vision Pro. Participants who saw Oculus also showed significantly higher

TABLE IV: Participants’ comprehension. The underlined percentages correspond to the correct answer. The red color highlights cases where the most common answer was incorrect. The green color highlights cases where the most common answer was correct. The Hol-Sys column corresponds to the Hololens system version of the question, Hol-App to HoloLens application, Oc-Sys to Oculus system, Oc-App to Oculus application, Vis-Sys to Vision Pro system, Vis-App to Vision Pro application.

Sensor	Category	Permission Comprehension Question	Options	Hol-Sys	Hol-App	Oc-Sys	Oc-App	Vis-Sys	Vis-App
Eye	Privacy	The system (application) requires your permission to access your eye-tracking data.	True	<u>92.5%</u>	<u>91.6%</u>	<u>96.6%</u>	<u>95.5%</u>	<u>87.1%</u>	<u>67.1%</u>
			False	2.8%	4.7%	1.1%	2.2%	9.4%	27.1%
			I Don't Know	4.7%	3.7%	2.3%	2.3%	3.5%	5.9%
		The system allows you to control which application has access to your eye-tracking data.	True	<u>87.9%</u>	N/A	<u>87.5%</u>	N/A	<u>55.3%</u>	N/A
			False	3.7%	N/A	3.4%	N/A	<u>31.8%</u>	N/A
			I Don't Know	8.4%	N/A	9.1%	N/A	13.0%	N/A
	The application can access your eye tracking data when running in the background.	True	N/A	<u>49.5%</u>	N/A	25.0%	N/A	29.4%	
		False	N/A	4.7%	N/A	20.5%	N/A	22.4%	
		I Don't Know	N/A	45.8%	N/A	<u>54.5%</u>	N/A	<u>48.2%</u>	
	The system (application) can transfer your eye-tracking data to an external device (e.g., a company server).	True	24.3%	<u>21.5%</u>	<u>23.9%</u>	<u>51.1%</u>	23.9%	24.7%	
		False	<u>26.2%</u>	17.8%	30.7%	9.1	<u>15.3%</u>	<u>32.9%</u>	
		I Don't Know	<u>49.5%</u>	<u>60.7%</u>	<u>39.8%</u>	<u>30.7%</u>	<u>56.5%</u>	<u>42.4%</u>	
	The system (application) can retain the unprocessed image of your eye.	True	<u>47.7%</u>	43.9%	35.2%	26.1%	43.5%	<u>47.1%</u>	
		False	11.2%	<u>7.5%</u>	<u>14.8%</u>	<u>15.9%</u>	<u>11.8%</u>	<u>16.5%</u>	
I Don't Know		41.1%	<u>48.6%</u>	<u>50.0%</u>	<u>58.0%</u>	<u>44.7%</u>	<u>36.5%</u>		
The system (application) only collects your final selection (instead of your eye movements) from the eye tracking data.	True	19.6%	21.5%	20.5%	23.9%	<u>18.8%</u>	<u>23.5%</u>		
	False	<u>26.2%</u>	<u>22.4%</u>	<u>25.0%</u>	<u>22.7%</u>	<u>18.8%</u>	<u>25.9%</u>		
	I Don't Know	<u>54.2%</u>	<u>56.1%</u>	<u>54.5%</u>	<u>53.4%</u>	<u>62.4%</u>	<u>50.6%</u>		
Utility	The system (application) can understand where your eyes look to indicate which virtual object to select.	True	<u>91.6%</u>	<u>72.9%</u>	<u>93.2%</u>	<u>93.2%</u>	<u>95.3%</u>	<u>81.2%</u>	
		False	0.9%	6.5%	3.4%	3.4%	1.2%	<u>8.2%</u>	
		I Don't Know	7.5%	20.6%	3.4%	3.4%	3.5%	10.6%	
	The system (application) can identify which real-world objects you are looking at.	True	<u>41.1%</u>	<u>34.6%</u>	<u>39.8%</u>	<u>44.3%</u>	<u>43.5%</u>	<u>40.0%</u>	
		False	19.6%	25.2%	34.1%	34.1%	24.7%	29.4%	
		I Don't Know	39.3%	<u>40.2%</u>	26.1%	21.6%	31.8%	30.6%	
The system (application) can simulate your eye movement for your virtual avatar.	True	<u>68.2%</u>	<u>59.8%</u>	<u>93.2%</u>	<u>95.5%</u>	<u>69.4%</u>	<u>67.1%</u>		
	False	8.4%	13.1%	3.4%	2.3%	5.9%	<u>8.2%</u>		
	I Don't Know	23.4%	33.3%	3.4%	2.3%	24.7%	24.7%		
The system (application) can authenticate your identity from the unique aspect of your eye (i.e., iris).	True	<u>84.1%</u>	<u>61.7%</u>	22.7%	25.0%	<u>85.9%</u>	<u>72.9%</u>		
	False	5.6%	9.3%	35.2%	34.1%	4.7%	15.3%		
	I Don't Know	10.3%	29.0%	<u>42.0%</u>	<u>40.9%</u>	9.4%	11.8%		
The system can adjust eye calibration for new users. The application can access user's eye calibration data	True	<u>79.4%</u>	<u>75.7%</u>	<u>86.4%</u>	<u>84.1%</u>	<u>76.5%</u>	<u>65.9%</u>		
	False	7.5%	4.7%	1.1%	4.5%	2.4%	14.1%		
	I Don't Know	13.1%	19.6%	12.5%	11.4%	21.2%	20.0%		
Hand	Privacy	The system (application) requires your permission to access your hand-tracking data.	True	<u>57.9%</u>	<u>71.0%</u>	<u>89.8%</u>	<u>55.7%</u>	<u>82.4%</u>	<u>96.5%</u>
			False	32.7%	21.5%	3.4%	34.1%	12.9%	1.2%
			I Don't Know	9.3%	7.5%	6.8%	10.2%	4.7%	2.4%
		The system allows you to control which application has access to your hand-tracking data.	True	39.3%	N/A	<u>40.9%</u>	N/A	<u>83.5%</u>	N/A
			False	<u>45.8%</u>	N/A	<u>38.6%</u>	N/A	5.9%	N/A
			I Don't Know	15.0%	N/A	20.5%	N/A	10.6%	N/A
	The application can access your hand-tracking data when running in the background.	True	N/A	<u>87.9%</u>	N/A	<u>55.7%</u>	N/A	37.6%	
		False	N/A	3.7%	N/A	9.1%	N/A	14.1%	
		I Don't Know	N/A	8.4%	N/A	35.2%	N/A	<u>48.2%</u>	
	The system (application) can transfer your hand-tracking data to an external device (e.g., a company server).	True	21.5%	<u>27.1%</u>	<u>35.2%</u>	34.1%	23.5%	<u>22.4%</u>	
		False	20.6%	15.9%	8.0%	11.4%	<u>12.9%</u>	12.9%	
		I Don't Know	<u>57.9%</u>	<u>57.0%</u>	<u>56.8%</u>	<u>54.5%</u>	<u>63.5%</u>	<u>64.7%</u>	
	The system (application) can retain the unprocessed image of your hand	True	<u>46.7%</u>	<u>48.6%</u>	48.9%	<u>55.7%</u>	49.4%	<u>61.2%</u>	
		False	12.1%	5.6%	9.1%	5.7%	10.6%	7.1%	
I Don't Know		41.1%	45.8%	42.0%	38.6%	40.0%	31.8%		
The system (application) only collects your final selection (instead of your hand movements) from the hand-tracking data.	True	12.1%	20.6%	19.3%	19.3%	24.7%	22.4%		
	False	27.1%	21.5%	27.3%	28.4%	9.4%	14.1%		
	I Don't Know	<u>60.7%</u>	<u>57.9%</u>	<u>53.4%</u>	<u>52.3%</u>	<u>65.9%</u>	<u>63.5%</u>		
Utility	The system (application) can understand your hand gesture to perform certain actions (e.g., select, scroll).	True	<u>92.5%</u>	<u>94.4%</u>	<u>100.0%</u>	<u>96.6%</u>	<u>95.3%</u>	<u>89.4%</u>	
		False	0.9%	0.9%	0.0%	1.1%	0.0%	1.2%	
		I Don't Know	7.5%	4.7%	0.0%	2.3%	4.7%	9.4%	
	The system (application) can identify which real-world objects you are holding.	True	<u>42.1%</u>	<u>43.9%</u>	28.4%	30.7%	<u>40.0%</u>	<u>32.9%</u>	
		False	18.7%	17.8%	28.4%	28.4%	22.4%	24.7%	
		I Don't Know	39.3%	38.3%	<u>43.2%</u>	<u>40.9%</u>	37.6%	<u>42.4%</u>	
The system (application) can simulate your hand movement for your virtual avatar.	True	<u>77.6%</u>	<u>74.8%</u>	<u>94.3%</u>	<u>93.2%</u>	<u>80.0%</u>	<u>72.9%</u>		
	False	3.7%	4.7%	1.1%	2.3%	2.4%	3.5%		
	I Don't Know	18.7%	20.6%	4.5%	4.5%	17.6%	23.5%		
The system (application) can authenticate your identity from the unique aspect of your hand (i.e., fingerprint).	True	24.3%	36.4%	22.7%	31.8%	<u>41.2%</u>	<u>42.4%</u>		
	False	31.8%	21.5%	34.1%	28.4%	21.2%	24.7%		
	I Don't Know	<u>43.9%</u>	<u>42.1%</u>	<u>43.2%</u>	<u>39.8%</u>	37.6%	32.9%		
The system (application) can measure the hand size of new users.	True	<u>42.1%</u>	<u>40.2%</u>	<u>59.1%</u>	<u>62.5%</u>	<u>64.7%</u>	<u>64.7%</u>		
	False	15.0%	12.1%	13.6%	14.8%	7.1%	4.7%		
	I Don't Know	<u>43.0%</u>	<u>47.7%</u>	27.3%	22.7%	28.2%	30.6%		

comprehension of privacy than participants who saw HoloLens ($z = 2.2008$, $p = .028$), but utility comprehension did not significantly differ ($z = 0.7158$, $p = .472$).

For hand-tracking, although participants who saw Oculus showed descriptively higher comprehension of the sensor utility, comprehension did not significantly differ from participants who saw Vision Pro ($z = 0.7762$, $p = .435$) or HoloLens ($z = 0.8611$, $p = .390$). Similarly, hand-tracking privacy comprehension among participants who saw Oculus was descriptively, but not significantly, higher than comprehension amongst participants who saw Vision Pro ($z = 0.3595$, $p = .719$) or HoloLens ($z = 0.1055$, $p = .912$).

C. RQ3: Specific Permission Comprehension

We next deeply investigate specific permission comprehension questions. We investigate the questions participants frequently answered incorrectly, indicated in a red color in Table IV. We identified six topics where misperceptions commonly occur. In some cases, participants underestimate privacy risks, and in other cases, they overestimate them (i.e., underestimate privacy protections). We also identified topics where participants showed good comprehension.

1) *Overestimating Access to Raw Data Retention:* Our results revealed a significant gap in participants' awareness regarding the system's or application's capability to retain unprocessed images of the eye or hand. When asked if the device could retain raw data, on average over 42% of respondents across all three platforms believed that the system could store unprocessed eye-tracking data, and over 48% believed the system could store raw hand-tracking data. In terms of the raw eye-tracking data, both HoloLens and Vision Pro retain identifiable iris data from users. While this information is encrypted on the device, it is important to ensure transparency in how these data are generated and how the raw data will be processed after iris patterns are generated. Oculus explicitly states that it does not store any raw eye-tracking data [10], yet only 14.8% of participants answered this question correctly. For hand-tracking, both HoloLens and Vision Pro stored processed hand-tracking, such as hand gestures for system interactions [33] and size and the shape of your hand [31]. Oculus again states that it does not store any raw hand-tracking data [15], but only 9.1% of participants answered this question correctly.

2) *Uninformed of Data Uploaded to System Servers:* Participants were also largely uninformed about these platforms' eye-tracking and hand-tracking data-sharing practices. At the system-level, we found that 30% of participants believed that Oculus does not share eye-tracking data with external servers, and another 40% were unsure. Oculus's privacy policy states that abstracted gaze data is sent to and *stored* on their servers and will be dissociated from individual accounts when they no longer need it [14]. Many participants who saw HoloLens (49.5%) and Vision Pro (56.5%) were also unsure whether eye-tracking data would be shared with an external server, though Vision Pro explicitly mentions that eye data will not be shared with Apple, and HoloLens states that it avoids

passing any identifiable information in their privacy statement. Similarly, we found that participants were largely uninformed about the hand-tracking data-sharing practice for HoloLens (57.9%), Oculus (56.8%), and Vision Pro (64.7%). Based on the privacy policy, we find that both HoloLens [33] and Vision Pro [31] stored abstracted hand-tracking information on the device, while Oculus shared the hand-tracking data with the Oculus server [15].

3) *Overestimating Permission Model for HoloLens and Vision Pro:* On a system level, we found that participants overestimated their ability to control the platform's access to their data through permissions. For example, 92.5% of participants believed they could control the system's access to eye-tracking data on HoloLens, and 87.1% believed the same for Vision Pro. For hand-tracking, 57.9% and 82.4% of participants held this assumption for HoloLens and Vision Pro, respectively. In reality, eye-tracking and hand-tracking for these systems are enabled by default. While both HoloLens and Vision Pro have adopted privacy-enhancing techniques to protect eye- and hand-tracking data, which are the primary source of interaction, our findings highlight a gap in user understanding of data control and practices on these platforms.

4) *Overestimation of Background Data Access for Oculus and Vision Pro:* Another common misunderstanding for Oculus and Vision Pro was the belief that applications could access eye-tracking or hand-tracking data in the background. From our experimentation, there was no direct API that allowed such background access. However, only a small percentage of participants answered correctly: 21.5% for eye-tracking and 11.6% for hand-tracking. For hand-tracking on HoloLens, the majority of participants (87.9%) answered correctly, likely due to the permission dialog clearly illustrating this capability.

5) *Overestimation Application Access To Calibration and Biometric Eye-tracking data:* At the application level, we found that participants overestimated the ability of applications to access their eye calibration data and biometric data (e.g., iris representation). For example, less than 5% of participants correctly identified that eye-tracking calibration data is not available to applications on HoloLens and Oculus. For platforms that support iris authentication, we found that on average, 67.3% of participants incorrectly believed that such information is accessible by applications.

6) *Uninformed about the Privacy Practice for Vision Pro:* Vision Pro acknowledges that even abstracted eye-tracking data could lead to serious privacy threats [31]. As a result, neither Apple nor third-party entities have access to these data. Only the final selection, rather than the eye movements, is available to the system and application. While Vision Pro arguably deploys the most privacy-preserving practices, we found that participants were largely uninformed. For example, when asked whether the system only collects the final selection, only 18.8% of participants answered correctly. Similarly, only 27.1% correctly understood that applications could not access eye-tracking data even with permission, and only 23.5% understood that only final selection is available to applications.

7) Comprehension for Eye- and Hand-Tracking Utility:

Finally, we highlight questions where participants, with no prior AR experience, showed good comprehension. When asking participants about the main utility for eye-tracking (simulating your eye movement), and for hand-tracking (simulating your hand movement), we found that the majority of the participants understand these utilities, especially when such utility is clearly illustrated. For example, in the case of the hand-tracking utility for the Oculus system, participants showed a 100% comprehension rate. The only exception is the application utility for eye-tracking on Vision Pro, which may be a result of our finding in Section V-C6.

D. RQ4: Factors Impacting User Decisions

Given our scope focusing on participants with no XR experience, we asked them to rate what information about the system and the app can help them feel more comfortable using the technology in the future. Post-experience, they selected three out of five factors we provided. Figure 6 illustrates the distribution of these factors, and we observe consistency in the factors across devices.

We observed that participants preferred information about who would have access to this data. Given the sensitivity of the collected biometric data, it is important to provide clear and comprehensive information on whether the system, external device (company server), or application developers will have access to their data. Additionally, clarity on how the data will be transmitted—whether it is encrypted, stored locally, or shared with remote servers—is crucial in building trust and comfort with the technology.

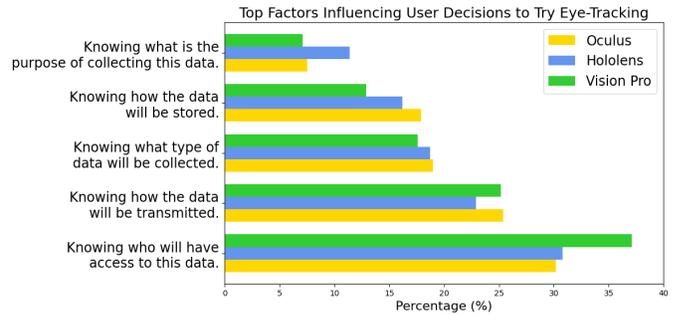
We were surprised to find less than 20% of participants regarded the type of collected data as a significant factor. This finding underscores a possible underestimation of the privacy risks associated with raw data access, and the need for more user education on the significance of raw data protection.

In addition, the distribution of factors considered important across hand-tracking and eye-tracking was highly consistent. This uniformity indicates that the factors that matter to users when considering the adoption of new technologies may remain consistent across various types of sensor data being collected. As AR technology advances, integrating more sophisticated sensors and collecting more data, we are hopeful that our findings will provide insights applicable to these emerging contexts as well.

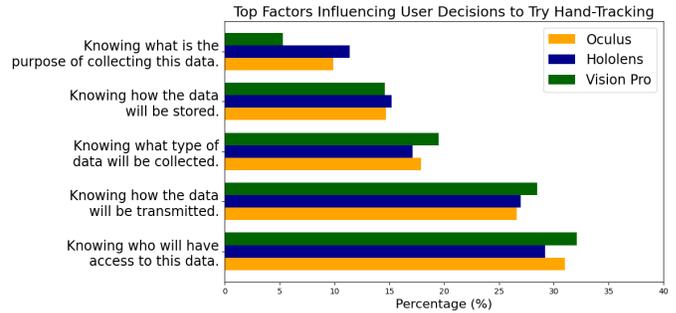
VI. DISCUSSION

User’s preference can be influenced by many factors, including the previous knowledge of these sensors, different data access models, dialog content, visualization, and UI flows. Informed by the results in Section V, we identified several key lessons from our work that could enhance user’s comfort and comprehension with implications for MR designers.

First, we found that effective communication about utility and privacy through permission UI flows enhances people’s comfort and willingness to use the technology, which aligns



(a) Decision Factors for Eye-tracking



(b) Decision Factors for Hand-tracking

Fig. 6: Decision factors breakdown for influencing user decisions to try eye-tracking and hand-tracking technologies on HoloLens, Oculus, and Vision Pro.

with previous studies [52], [54], [99]. For example, the permission flow on Oculus provides more detailed descriptions, compared to HoloLens and Vision Pro, regarding the utility and privacy implications of eye- and hand-tracking sensors. Consequently, people who interacted with the Oculus interface were better informed, which not only aligned with how informed they *felt*, but also increased their comfort compared to the other platforms. Our findings in Section V-C suggest the necessity of including relevant descriptions to enhance topics where users tend to underestimate the system’s privacy protections, such as preventing the retention of raw data.

Suggestion 1: AR platforms and developers should provide clear communication on potential utility and privacy to enhance user comfort and comprehension.

Second, our findings in Section III illustrate the different approaches AR platforms have taken in handling users’ eye-tracking data. While HoloLens and Oculus both took active steps to protect users’ privacy by only providing abstracted eye-tracking data, recent studies have shown that even abstracted eye-tracking data can contain significant privacy risks, such as revealing user intention [44], [51], psychological state [90], age [37], and cultural background [82]. Hence, we encourage these platforms to explore potential privacy-preserving mechanisms, including limiting system and application’s access [31] or adding stronger privacy guarantees over the abstracted eye-tracking data stream [70], [73], [91]. However, we also noticed that while Vision Pro adopted

stronger privacy-preserving techniques (providing to apps only final UI selections the system derives from eye-tracking data), people often failed to fully comprehend their implications (see Section V-C6). Apple may be missing an opportunity to enhance user comfort and understanding by clearly explaining the deployed protections.

Suggestion 2: Given the potential privacy risks with even abstracted eye-tracking data, we encourage platforms such as Oculus and HoloLens to provide stronger privacy protection. We also advocate for better communication with users if such practice is adopted.

Third, we found that while many participants considered data transmission information important for both sensors, many were inadequately informed about this factor. For example, around 70% of the participants were unaware that Oculus shares eye-tracking data to its own external server (and over 90% for hand-tracking). Despite Oculus outlining this in their privacy policy, given the low likelihood of users reading privacy policies [60], [63], [85], such information is not effectively communicated. We argue that it is essential to implement opt-in/opt-out features, allowing users to control their data-sharing preferences. In addition, the eye-tracking data retention period needs a clearer definition than “deletion when no longer needed”. Stepping back, platforms should also consider whether this data needs to be shared with external servers at all, and at what granularity and for which purposes.

Suggestion 3: For platforms that do upload data, such as Oculus, we suggest: (a) implementing an opt-in/opt-out feature for users to choose whether they wish to share eye-tracking data with external servers, and (b) providing a transparent explanation for this data collection, including the retention period, in the permission flow.

Fourth, our findings in Section III suggest that HoloLens and Oculus grant applications automatic access to users’ hand-tracking data, but a minority of participants understood this. Recognizing that hand-tracking is the main interaction modality and cannot be realistically opted out of entirely, we recommend that HoloLens and Oculus still provide finer-grained permission to limit applications’ access to certain hand-tracking data. For example, precise estimation of hand skeleton data could be limited, given its potential privacy implications for inferring sensitive attributes [71], [78].

Suggestion 4: AR platforms should implement fine-grained permissions for hand-tracking to provide users more control over their data, e.g., by restricting applications’ access to specific types of hand-tracking data.

Finally, it is important to consider the trade-off between privacy and usability in our recommendations. For example, deploying fine-grained hand-tracking permissions might put an extra burden on users. However, our results in Section V-C3 suggest that the majority of participants expect this control from HoloLens (71.0%) and Oculus (82.4%). In addition, results in Section V-A1 suggested that participants who saw HoloLens felt less comfortable with sharing hand-tracking data compared to Oculus and Vision Pro, possibly due to the lack of hand-tracking permissions. Similarly, if AR systems clearly illustrate their privacy mechanisms for eye- or hand-tracking data, this transparency might deter new users who are concerned about potential privacy issues. Although AR technologies have grown significantly over the past few years, with initial Vision Pro sales estimate of 200,000 units in 2024 [30], Quest Pro sales estimate of 100,000 units [22]–[25], [32], and HoloLens 2 sales estimate of 300,000 units [6] since release, they are still in the early stages of mass adoption. To position these technologies for broader acceptance, it is crucial to enhance users’ trust through effective privacy mechanisms [40], preparing the mainstream market to bridge the adoption chasm [76]. We encourage future research to further explore this direction.

Suggestion 5: Proper privacy-enhancing techniques can better prepare AR technologies for future widespread adoption.

VII. CONCLUSION

As AR technologies advance, novel privacy concerns also emerge. We sought to explore how three existing AR headsets — Meta’s Oculus Pro, Microsoft’s HoloLens 2, and Apple’s Vision Pro — navigate permissions for eye- and hand-tracking, and the extent to which users feel comfortable and informed about these sensors. We find that people’s experiences with and comprehension of permissions are affected by both the different design choices across devices and the sensors themselves. Based on our findings, we suggest how future AR platforms can design permissions that effectively communicate information that is particularly important and that often goes misunderstood by end users.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their valuable feedback. This work was supported in part by the National Science Foundation under Award 2205171, 2114230, and 2207019, as well as by awards from Cisco, Google, and Qualcomm.

REFERENCES

- [1] Microsoft hololens 2. <https://www.microsoft.com/en-us/hololens>, 2019.
- [2] Eye tracking on meta quest pro. <https://www.meta.com/help/quest/articles/getting-started/getting-started-with-quest-pro/eye-tracking/>, 2022.
- [3] Hand tracking gestures - hololens. <https://learn.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/mrtk2/features/input/gestures?view=mrtkunity-2022-05>, 2022.

- [4] Hand tracking — mrtk2. <https://learn.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/mrtk2/features/input/hand-tracking?view=mrtkunity-2022-05>, 2022.
- [5] Locatable camera overview. <https://learn.microsoft.com/en-us/windows/mixed-reality/develop/advanced-concepts/locatable-camera-overview>, 2022.
- [6] Microsoft has sold 300,000 hololens units according to analysts. <https://www.thurrott.com/mobile/275228/microsoft-hololens-300000-units-sold>, 2022.
- [7] Adopting best practices for privacy and user preferences - vision pro. <https://developer.apple.com/documentation/visionos/adopting-best-practices-for-privacy>, 2023.
- [8] Apple vision pro. <https://www.apple.com/apple-vision-pro/>, 2023.
- [9] Apple vision pro - press release. <https://www.microsoft.com/en-us/hololens/hardware>, 2023.
- [10] Building Eye Tracking on Meta Quest Pro Responsibly . https://scontent-sea1-1.xx.fbcdn.net/v/t39.8562-6/312898144_1269308143870038_8244941952542354869_n.pdf?_nc_cat=111&ccb=1-7&_nc_sid=b8d81d&_nc_ohc=t8tptL9ReZ0AX-ZulfQ&_nc_ht=scontent-sea1-1.xx&oh=00_AfAiD1zsX1vS-jlTgFodXOrETENsUSz9VXDDh4aTDIxDPQ&oe=654F5076, 2023.
- [11] Enable hand tracking on oculus. <https://developer.oculus.com/documentation/native/android/mobile-hand-tracking/>, 2023.
- [12] Extended eye tracking in native hololens engine. <https://learn.microsoft.com/en-us/windows/mixed-reality/develop/native/extended-eye-tracking-native>, 2023.
- [13] Eye tracked foveated rendering. <https://developer.oculus.com/documentation/unity/unity-eye-tracked-foveated-rendering/>, 2023.
- [14] Eye tracking privacy notice for oculus. <https://www.meta.com/help/quest/articles/accounts/privacy-information-and-settings/eye-tracking-privacy-notice/>, 2023.
- [15] Hand and body privacy notice - oculus. <https://www.meta.com/help/quest/articles/accounts/privacy-information-and-settings/hand-tracking-privacy-notice/>, 2023.
- [16] Hand tracking gestures - oculus. <https://www.meta.com/help/quest/articles/headsets-and-accessories/controllers-and-hand-tracking/hand-tracking/>, 2023.
- [17] Hand tracking gestures detection - vision pro. https://developer.apple.com/documentation/vision/detecting_hand_poses_with_vision, 2023.
- [18] Hand tracking on meta quest. <https://www.meta.com/help/quest/articles/headsets-and-accessories/controllers-and-hand-tracking/hand-tracking/>, 2023.
- [19] Handtrackingprovider from visionos. <https://developer.apple.com/documentation/arkit/handrackingprovider>, 2023.
- [20] Hololens 2 native eye-tracking api. <https://learn.microsoft.com/en-us/uwp/api/windows.perception.people.eyespose?view=wint-22621>, 2023.
- [21] Hololens 2 technical specifications. <https://www.apple.com/newsroom/2023/06/introducing-apple-vision-pro/>, 2023.
- [22] How many vr headsets did meta sell in q1 2023? <https://arinsider.co/2023/05/01/how-many-vr-headsets-did-meta-sell-in-q1/>, 2023.
- [23] How many vr headsets did meta sell in q2 2023? <https://arinsider.co/2023/07/31/how-many-vr-headsets-did-meta-sell-in-q2/>, 2023.
- [24] How many vr headsets did meta sell in q3 2023? <https://arinsider.co/2023/10/30/how-many-vr-headsets-did-meta-sell-in-q3-2/>, 2023.
- [25] How many vr headsets did meta sell in q4 2022? <https://arinsider.co/2023/02/06/how-many-vr-headsets-did-meta-sell-in-q4/>, 2023.
- [26] Improve visual quality and comfort - hololens. <https://learn.microsoft.com/en-us/hololens/hololens-calibration>, 2023.
- [27] Manage user identity and login for hololens. <https://learn.microsoft.com/en-us/hololens/hololens-identity>, 2023.
- [28] Meta Quest Pro . , 2023.
- [29] Permissions required to use the dynamics 365 guides hololens app. <https://learn.microsoft.com/en-us/dynamics365/mixed-reality/guides/hololens-permissions>, 2023.
- [30] Apple has sold approximately 200,000 vision pro headsets. <https://www.macrumors.com/2024/01/29/apple-vision-pro-headset-sales/>, 2024.
- [31] Apple vision pro privacy overview. https://www.apple.com/privacy/docs/Apple_Vision_Pro_Privacy_Overview.pdf, 2024.
- [32] How many vr headsets did meta sell in q4 2023? <https://arinsider.co/2024/02/12/how-many-headsets-did-meta-sell-in-q4/>, 2024.
- [33] Microsoft privacy statement. <https://privacy.microsoft.com/en-us/privacy-statement>, 2024.
- [34] Openxr. <https://registry.khronos.org/OpenXR/specs/1.0/html/xrspec.html>, 2024.
- [35] Swift UI. <https://developer.apple.com/xcode/swiftui/>, 2024.
- [36] UIKit. <https://developer.apple.com/documentation/uikit>.
- [37] Alper Aık, Adjal Sarwary, Rafael Schultze-Kraft, Selim Onat, and Peter Konig. Developmental changes in natural viewing behavior: bottom-up and top-down differences between children, young adults and older adults. *Frontiers in psychology*, 1:7198, 2010.
- [38] Surin Ahn, Maria Gorlatova, Parinaz Naghizadeh, Mung Chiang, and Prateek Mittal. Adaptive fog-based output security for augmented reality. In *Proceedings of the Morning Workshop on Virtual Reality and Augmented Reality Network*, pages 1–6, 2018.
- [39] Benjamin Bach, Ronell Sicat, Johanna Beyer, Maxime Cordeil, and Hanspeter Pfister. The hologram in my hand: How effective is interactive exploration of 3d visualizations in immersive tangible augmented reality? *IEEE transactions on visualization and computer graphics*, 24(1):457–467, 2017.
- [40] France Belanger and Robert E Crossler. Privacy in the digital age: a review of information privacy research in information systems. *MIS quarterly*, pages 1017–1041, 2011.
- [41] Kevin Benton, L Jean Camp, and Vaibhav Garg. Studying the effectiveness of android application permissions requests. In *2013 IEEE international conference on pervasive computing and communications workshops (PERCOM Workshops)*, pages 291–296. IEEE, 2013.
- [42] Mara Jose Blanca, Jaume Arnau, Javier Garca-Castro, Rafael Alarcon, and Roser Bono. Non-normal data in repeated measures anova: impact on type i error and power. *Psicothema*, pages 21–29, 2023.
- [43] Jonas Blattgerste, Patrick Renner, and Thies Pfeiffer. Advantages of eye-gaze over head-gaze-based selection in virtual and augmented reality under varying field of views. In *Proceedings of the Workshop on Communication by Gaze Interaction*, pages 1–9, 2018.
- [44] Sylvain Castagnos, Nicolas Jones, and Pearl Pu. Eye-tracking product recommenders’ usage. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 29–36, 2010.
- [45] Kaiming Cheng, Arkaprabha Bhattacharya, Michelle Lin, Jaewook Lee, Aroosh Kumar, Jeffery F Tian, Tadayoshi Kohno, and Franziska Roesner. When the user is inside the user interface: An empirical study of ui security properties in augmented reality. In *USENIX Security*, 2024.
- [46] Kaiming Cheng, Jeffery F Tian, Tadayoshi Kohno, and Franziska Roesner. Exploring user reactions and mental models towards perceptual manipulation attacks in mixed reality. In *USENIX Security*, volume 18, 2023.
- [47] James Crowley, Franois Berard, Joelle Coutaz, et al. Finger tracking as an input device for augmented reality. In *International Workshop on Gesture and Face Recognition*, pages 195–200, 1995.
- [48] Jaybie A De Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. Security and privacy approaches in mixed reality: A literature survey. *ACM Computing Surveys (CSUR)*, 52(6):1–37, 2019.
- [49] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2377–2386, 2014.
- [50] Andrew T Duchowski. Gaze-based interaction: A 30 year retrospective. *Computers & Graphics*, 73:59–69, 2018.
- [51] Sukru Eraslan, Yeliz Yesilada, and Simon Harper. Scanpath trend analysis on web pages: Clustering eye tracking scanpaths. *ACM Transactions on the Web (TWEB)*, 10(4):1–35, 2016.
- [52] Adrienne Porter Felt, Serge Egelman, Matthew Finifter, Devdatta Akhawe, David A Wagner, et al. How to ask for permission. *HotSec*, 12:7–7, 2012.
- [53] Adrienne Porter Felt, Serge Egelman, and David Wagner. I’ve got 99 problems, but vibration ain’t one: a survey of smartphone users’ concerns. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*, pages 33–44, 2012.
- [54] Adrienne Porter Felt, Kate Greenwood, and David Wagner. The effectiveness of application permissions. In *2nd USENIX Conference on Web Application Development (WebApps 11)*, 2011.

- [55] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–14, 2012.
- [56] Adrienne Porter Felt, Robert W Reeder, Hazim Almuhammedi, and Sunny Consolvo. Experimenting at scale with google chrome’s ssl warning. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2667–2670, 2014.
- [57] Andrea Gallardo, Chris Choy, Jaideep Juneja, Efe Bozkir, Camille Cobb, Lujo Bauer, and Lorrie Cranor. Speculative privacy concerns about ar glasses data collection. *Proceedings on Privacy Enhancing Technologies*, 4:416–435, 2023.
- [58] Jassim Happa, Mashhuda Glencross, and Anthony Steed. Cyber security threats and challenges in collaborative mixed-reality. *Frontiers in ICT*, 6:5, 2019.
- [59] David Harborth and Alisa Frik. Evaluating and redefining smartphone permissions with contextualized justifications for mobile augmented reality apps. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, pages 513–534, 2021.
- [60] Duha Ibdah, Nada Lachtar, Satya Meenakshi Raparathi, and Anys Bacha. “why should i read the privacy policy, i just need the service”: A study on attitudes and perceptions toward privacy policies. *IEEE access*, 9:166465–166487, 2021.
- [61] Suman Jana, David Molnar, Alexander Moshchuk, Alan Dunn, Benjamin Livshits, Helen J Wang, and Eyal Ofek. Enabling fine-grained permissions for augmented reality applications with recognizers. In *22nd USENIX Security Symposium*, pages 415–430, 2013.
- [62] Suman Jana, Arvind Narayanan, and Vitaly Shmatikov. A scanner darkly: Protecting user privacy from perceptual applications. In *IEEE Symposium on Security and Privacy*, pages 349–363, 2013.
- [63] Farzaneh Karegar, John Sören Pettersson, and Simone Fischer-Hübner. The dilemma of user engagement in privacy notices: Effects of interaction modes and habituation on user attention. *ACM Transactions on Privacy and Security (TOPS)*, 23(1):1–38, 2020.
- [64] George Alex Koulieris, Kaan Akşit, Michael Stengel, Rafał K Mantiuk, Katerina Mania, and Christian Richardt. Near-eye display and tracking technologies for virtual and augmented reality. In *Computer Graphics Forum*, volume 38, pages 493–519. Wiley Online Library, 2019.
- [65] Gregory Kramida. Resolving the vengeance-accommodation conflict in head-mounted displays. *IEEE transactions on visualization and computer graphics*, 22(7):1912–1931, 2015.
- [66] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. What does your gaze reveal about you? on the privacy implications of eye tracking. *Privacy and Identity Management. Data for Better Living: AI and Privacy: 14th IFIP WG 9.2, 9.6/11.7, 11.6/SIG 9.2. 2 International Summer School, Windisch, Switzerland, August 19–23, 2019, Revised Selected Papers 14*, pages 226–241, 2020.
- [67] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. Securing augmented reality output. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 320–337. IEEE, 2017.
- [68] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. Towards security and privacy for multi-user augmented reality: Foundations with end users. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 392–408. IEEE, 2018.
- [69] Hyunjoon Lee, Jiyeon Lee, Daejun Kim, Suman Jana, Insik Shin, and Soeul Son. AdCube: WebVR Ad Fraud and Practical Confinement of Third-Party Ads. In *USENIX Security Symposium*, pages 2543–2560, 2021.
- [70] Jingjie Li, Amrita Roy Chowdhury, Kassem Fawaz, and Younghyun Kim. {Kaleido}:{Real-Time} privacy control for {Eye-Tracking} systems. In *30th USENIX security symposium (USENIX security 21)*, pages 1793–1810, 2021.
- [71] Jonathan Liebers, Sascha Brockel, Uwe Gruenefeld, and Stefan Schneegass. Identifying users by their hand tracking data in augmented and virtual reality. *International Journal of Human-Computer Interaction*, pages 1–16, 2022.
- [72] Daniel J Liebling and Sören Preibusch. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, pages 1169–1177, 2014.
- [73] Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. Differential privacy for eye-tracking data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–10, 2019.
- [74] Tongbo Luo, Xing Jin, Ajai Ananthanarayanan, and Wenliang Du. Touchjacking attacks on web in android, ios, and windows phone. In *International Symposium on Foundations and Practice of Security*, pages 227–243. Springer, 2012.
- [75] Shahzad Malik, Chris McDonald, and Gerhard Roth. Hand tracking for interactive pattern-based augmented reality. In *Proceedings. International Symposium on Mixed and Augmented Reality*, pages 117–126. IEEE, 2002.
- [76] Geoffrey A Moore and Regis McKenna. Crossing the chasm. 1999.
- [77] Alexios Mylonas, Marianthi Theoharidou, and Dimitris Gritzalis. Assessing privacy risks in android: A user-centric approach. In *Risk Assessment and Risk-Driven Testing: First International Workshop, RISK 2013, Held in Conjunction with ICTSS 2013, Istanbul, Turkey, November 12, 2013. Revised Selected Papers I*, pages 21–37. Springer, 2014.
- [78] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F O’Brien, Louis Rosenbergs, and Dawn Song. Unique identification of 50,000+ virtual reality users from head & hand motion data. *arXiv preprint arXiv:2302.08927*, 2023.
- [79] Joseph O’Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. Privacy-enhancing technology and everyday augmented reality: Understanding bystanders’ varying needs for awareness and consent. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4):1–35, 2023.
- [80] Ken Pfeuffer, Matthias J Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. Behavioural biometrics in vr: Identifying people from body motion and relations in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2019.
- [81] Philipp A Rauschnabel, Jun He, and Young K Ro. Antecedents to the adoption of augmented reality smart glasses: A closer look at privacy risks. *Journal of Business Research*, 92:374–384, 2018.
- [82] Keith Rayner, Monica S Castelano, and Jimmian Yang. Eye movements when looking at unusual/weird scenes: Are there cultural differences? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(1):254, 2009.
- [83] Franziska Roesner and Tadayoshi Kohno. Security and privacy for augmented reality: Our 10-year retrospective. In *VR4Sec: 1st International Workshop on Security for XR and XR for Security*, 2021.
- [84] Franziska Roesner, Tadayoshi Kohno, and David Molnar. Security and privacy for augmented reality systems. *Communications of the ACM*, 57(4):88–96, 2014.
- [85] Manuel Rudolph, Denis Feth, and Svenja Polst. Why users ignore privacy policies—a survey and intention model for explaining user privacy behavior. In *Human-Computer Interaction. Theories, Methods, and Human Issues: 20th International Conference, HCI International 2018, Las Vegas, NV, USA, July 15–20, 2018, Proceedings, Part I 20*, pages 587–598. Springer, 2018.
- [86] Bingyu Shen, Lili Wei, Chengcheng Xiang, Yudong Wu, Mingyao Shen, Yuanyuan Zhou, and Xinxin Jin. Can systems explain permissions better? understanding users’ misperceptions under smartphone runtime permission model. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 751–768. USENIX Association, August 2021.
- [87] Yan Shen, Soh-Khim Ong, and Andrew YC Nee. Vision-based hand interaction in augmented reality environment. *Intl. Journal of Human-Computer Interaction*, 27(6):523–544, 2011.
- [88] Gulshan Shrivastava, Prabhat Kumar, Deepak Gupta, and Joel JPC Rodrigues. Privacy issues of android application permissions: A literature review. *Transactions on Emerging Telecommunications Technologies*, 31(12):e3773, 2020.
- [89] Carter Slocum, Yicheng Zhang, Nael Abu-Ghazaleh, and Jiasi Chen. Going through the motions: AR/VR keylogging from user head motions. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 159–174, 2023.
- [90] Paula C Stacey, Stephanie Walker, and Jean DM Underwood. Face processing and familiarity: Evidence from eye-movement data. *British Journal of Psychology*, 96(4):407–422, 2005.
- [91] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–9, 2019.

- [92] William Steptoe, Anthony Steed, Aitor Rovira, and John Rae. Lie tracking: social presence, truth and deception in avatar-mediated telecommunication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1039–1048, 2010.
- [93] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX security symposium (USENIX security 22)*, pages 3789–3806, 2022.
- [94] John Vilk, David Molnar, Benjamin Livshits, Eyal Ofek, Chris Rossbach, Alexander Moshchuk, Helen J Wang, and Ran Gal. SurroundWeb: Mitigating privacy concerns in a 3D web browser. In *IEEE Symposium on Security and Privacy*, pages 431–446, 2015.
- [95] Vinoba Vinayagamoorthy, Andrea Brogni, Marco Gillies, Mel Slater, and Anthony Steed. An investigation of presence response across variations in visual realism. In *The 7th Annual International Presence Workshop*, pages 148–155, 2004.
- [96] Frederike Wenzlaff, Peer Briken, and Arne Dekker. Video-based eye tracking in sex research: a systematic literature review. *The Journal of Sex Research*, 53(8):1008–1019, 2016.
- [97] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. Android permissions remystified: A field study on contextual integrity. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 499–514, 2015.
- [98] Yi Wu, Cong Shi, Tianfang Zhang, Payton Walker, Jian Liu, Nitesh Saxena, and Yingying Chen. Privacy leakage via unrestricted motion-position sensors in the age of virtual reality: A study of snooping typed input on virtual keyboards. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 3382–3398. IEEE Computer Society, 2023.
- [99] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. Towards effective differential privacy communication for users’ data sharing decision and comprehension. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 392–410. IEEE, 2020.
- [100] Songhua Xu, Hao Jiang, and Francis CM Lau. Personalized online document, image and video recommendation via commodity eye-tracking. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 83–90, 2008.

VIII. RECRUITMENT & SURVEY

In our study, we asked the same set of questions across three different devices. The only difference was the corresponding information about the permission UI, such as permission dialog screenshots and whether the device prompted for permission. We provided alt-text for all screenshots in our survey.

[Recruiting Message] In this study, we are hoping to evaluate the permission-granting process of current Augmented/Mixed Reality Headsets (Apple Vision Pro, Hololens 2, and Oculus Quest Pro). You will be asked to complete a questionnaire which will take around 13 minutes. We are looking for participants who have little or no experience with Augmented/Mixed Reality headsets. When taking the survey, simply answer the questions as honestly as you can. Thank you for your interest in this research.

[Consent Form] Thank you for taking the survey! We are a group of researchers from the University of Washington, and we are hoping to evaluate the permission-granting process of current Augmented/Mixed Reality Headsets (Apple Vision Pro, Hololens 2, and Oculus Quest Pro). You will be asked to complete a questionnaire which will take around 10 minutes. This study was reviewed by the UW Institutional Review Board (IRB) and deemed exempt because it involves no more than minimal risk and meets other criteria. Your

responses to this survey will be anonymized. Data from this survey will be stored securely and kept confidential. Your participation in this survey is voluntary, and you may withdraw anytime. If you have questions about this study, please contact Kaiming Cheng (Ph.D. candidate at UW) at kaimingc@cs.washington.edu. You may also contact the UW Human Subjects Division (HSD), which manages IRB review, at hsdinfo@uw.edu. Thank you for taking our survey!

[Filtering] Do you consent to participate in this study?

- (i) I am at least 18 years old, I have read and understood this consent form, and I agree to participate in this online research study.
- (ii) I do not wish to participate in this study.

[Context] Welcome to the study! We are investigating user perceptions and comfort with the permission-granting process in Augmented and Mixed Reality technologies. Augmented Reality/Mixed Reality (AR/MR) is a technology that overlays digital information onto a user’s view of the real world. One common Augmented/Mixed Reality device is a Head Mounted Display, or a headset. AR/MR headsets come in various forms - from looking like regular glasses to looking more like helmets. For example, here are some existing AR/MR headsets on the market today (Figure 7)

[Tech Background] Do you have a background in technology through education or professional experience?

- (i) Yes
- (ii) No

[AR Familiarity] Have you heard of Augmented Reality/Mixed Reality (AR/MR) before this study?

- (i) Yes
- (ii) No

[AR Experience] What is your experience level with Augmented Reality/Mixed Reality (AR/MR) headsets?

- (i) I have never used any AR/MR headset.
- (ii) I have used an AR/MR headset a few times.
- (iii) I am an active user of AR/MR headsets.

[AR Headset Usage] If you have used any of the following AR/MR headsets: Microsoft Hololens 2, Apple Vision Pro, or Meta Quest 3, please select those devices below.

- (i) Microsoft Hololens 2
- (ii) Apple Vision Pro
- (iii) Meta Quest 3
- (iv) I have not used any of the above devices.

A. Survey for Eye-tracking on Oculus

[Introduction] Augmented and Mixed Reality headsets have a variety of sensors recording data while the headset is in use. Users of these headsets typically view permission dialogs to let you allow or deny this request to access your data for different sensors. In this survey, we will present permission dialogs for two different types of sensors and ask for your impressions of each set of dialogs. When you continue, you will see the first sensor.



Fig. 7: Images of AR headsets (Meta's Quest Pro, Microsoft's HoloLens 2, and Apple's Vision Pro)

[Instruction] Suppose you want to use an AR/MR headset with an eye-tracking feature. Below is what you see in the process of granting permission for eye-tracking. We would like to ask you about your comfort levels and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for eye tracking. You can enable eye-tracking permission, pause eye tracking, and control eye calibration data for the system in this dialog from the system setting. After you toggle the button, the following dialog appears (Figure 8):

After you enable the eye-tracking feature, you will be asked to perform a calibration process. You can control which application has access to your eye-tracking data in the system setting (Figure 9):

[Q1]: I feel informed about the utility of this permission. (5pt Likert scale from "Strong disagree" to "Strong agree")

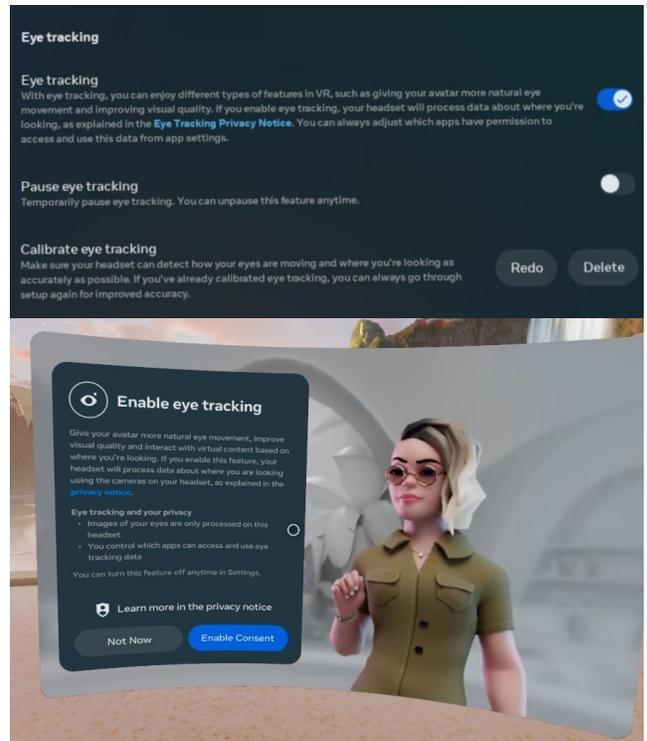


Fig. 8: System-level eye-tracking permission dialog (Oculus)

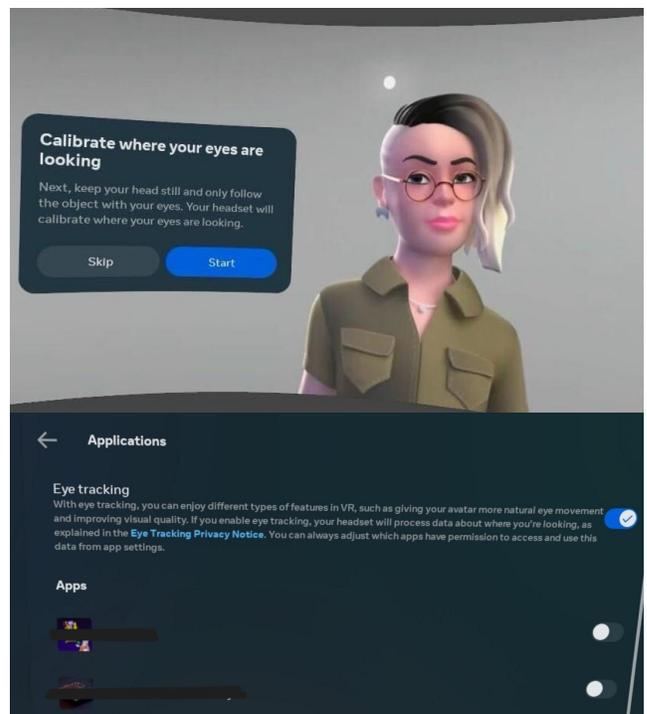


Fig. 9: Eye-tracking calibration and app permission control (Oculus)

[Q2]: I feel informed about the associated privacy risk of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q3]: I feel confident that this AR/MR system will securely store my eye-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q4]: I know exactly what type data will be collected, how it will be used, and who will have access to it based on the information presented in the above permission screenshots. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q5]: I feel comfortable using the device knowing the level of access it has to my eye tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Instruction] Now, we are interested in the degree to which you understand what the system (i.e., the headset) can do with your data once you grant permission. Answer the following true or false questions regarding the **sensor capability**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q6]: The system can understand where your eyes look to indicate which virtual object to select. 1. True 2. False 3. I don't know

[Q7]: The system can identify which real-world objects you are looking at. 1. True 2. False 3. I don't know

[Q8]: The system can simulate your eye movement for your virtual avatar. 1. True 2. False 3. I don't know

[Q9]: The system can authenticate your identity from the unique aspect of your eye (i.e., iris). 1. True 2. False 3. I don't know

[Q10]: The system can adjust eye calibration for new users. 1. True 2. False 3. I don't know

[Instruction] Answer the following true or false questions regarding the **sensor privacy**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q11]: The system requires your permission to access your eye tracking data. 1. True 2. False 3. I don't know

[Q12]: The system allows you to control which application has access to your eye tracking. 1. True 2. False 3. I don't know

[Q13]: The system can transfer your eye tracking data to an external device (e.g., a company server). 1. True 2. False 3. I

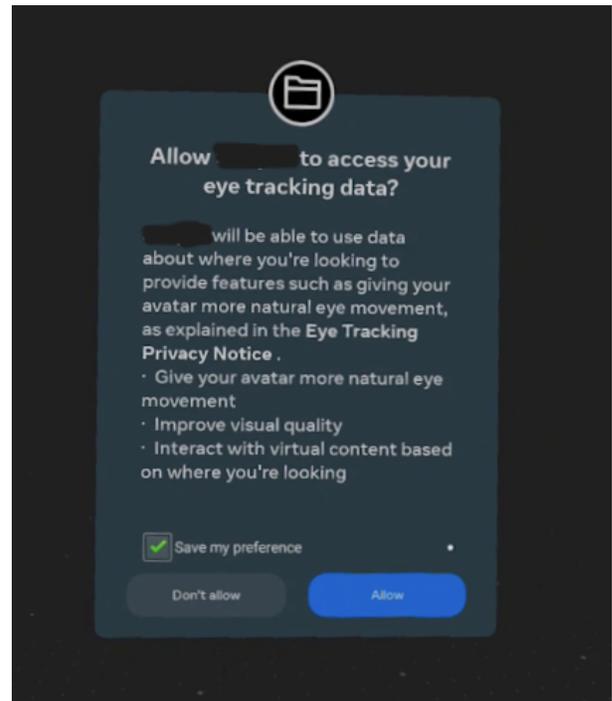


Fig. 10: App-level eye-tracking permission dialog (Oculus)

don't know

[Q14]: The system can retain the unprocessed image of your eye on the AR/MR headset. 1. True 2. False 3. I don't know

[Q15]: The system only collects your final selection (instead of your eye movements) from the eye tracking data. 1. True 2. False 3. I don't know

[Instruction] Now, you open an app on the headset, which has its own app-level permission settings for eye tracking. The following app dialog appears after you open the application for the first time (Figure 10):

[Q16]: I feel informed about the utility of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q17]: I feel informed about the associated privacy risk of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q18]: I feel confident that this AR/MR application will securely store my eye-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q19]: I know exactly what type data will be collected, how it will be used, and who will have access to it based on the information presented in the above permission screenshots. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q20]: I feel comfortable using this AR/MR application knowing the level of access it has to my eye-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Instruction] Now, we are interested in the degree to which you understand what the application can do with your data once you grant permission. Answer the following true or false questions regarding the **sensor capability**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q21]: The application can understand where your eyes look to indicate which virtual object to select. 1. True 2. False 3. I don’t know

[Q22]: The application can identify which real-world objects you are looking at. 1. True 2. False 3. I don’t know

[Q23]: The application can simulate your eye movement for your virtual avatar. 1. True 2. False 3. I don’t know

[Q24]: The application can authenticate your identity from the unique aspect of your eye (i.e., iris). 1. True 2. False 3. I don’t know

[Q25]: The application can access user’s eye calibration data (e.g., eye position) provided by the system. 1. True 2. False 3. I don’t know

[Instruction] Answer the following true or false questions regarding the sensor privacy. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q26]: The application requires your permission to access your eye tracking data. 1. True 2. False 3. I don’t know

[Q27]: The application can access your eye tracking data when running in the background. 1. True 2. False 3. I don’t know

[Q28]: The application can transfer your eye tracking data to an external device (e.g., a company server). 1. True 2. False 3. I don’t know

[Q29]: The application can retain the unprocessed image of your eye within the application. 1. True 2. False 3. I don’t know

[Q30]: The application only collects your final selection (instead of your eye movements) from the eye tracking data. 1. True 2. False 3. I don’t know

[Instruction] Now that you have seen the permission settings for both the overall system and the app, we want to understand what information about both the system and the app can help

you feel more comfortable using the technology in the future.

[Instruction] Please drag and drop the top three most important items from the list below that can influence your decision to use this technology in the future. (Don’t worry about the ordering within the box)

[Item 1]: Knowing who will have access to this data. [Example includes: permission request; background access, control which app has access to your data].

[Item 2]: Knowing how will the data be stored. [Example includes: Delete after use, stores eye tracking data by default; provide options to delete your data.]

[Item 3]: Knowing how will the data be transmitted. [Example includes: keep your data only on device; transfer your data to an external device]

[Item 4]: Knowing what type of data will be collected. [Example includes: eye movement data (how long you look); eye gaze data (where you look); final selection (where you indicate); unique aspect of your eye (iris).]

[Item 5]: Knowing what is the purpose of collecting this data. [Example includes: indicate selection; generate virtual avatar; identity authentication]

B. Survey for Hand-tracking on Oculus

[Instruction] Suppose you want to use an AR/MR application with a hand-tracking feature. Below is what you see in the process of granting permission for hand tracking. We would like to ask you about your comfort levels and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for hand tracking. You can enable hand-tracking permission for the system in this dialog from device permission in the system setting. After you toggle the button, the following dialog appears (Figure 11):

[Hand-tracking tutorial] After you enable the hand-tracking feature, the system will present tutorials on how to interact with the virtual content using your hand (Figure 12):

[Q31]: I feel informed about the utility of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q32]: I feel informed about the associated privacy risk of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q33]: I feel confident that this AR/MR system will securely store my hand-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q34]: I know exactly what type data will be collected, how it will be used, and who will have access to it based on the information presented in the above permission screenshots. (5pt Likert scale from “Strong disagree” to “Strong agree”)

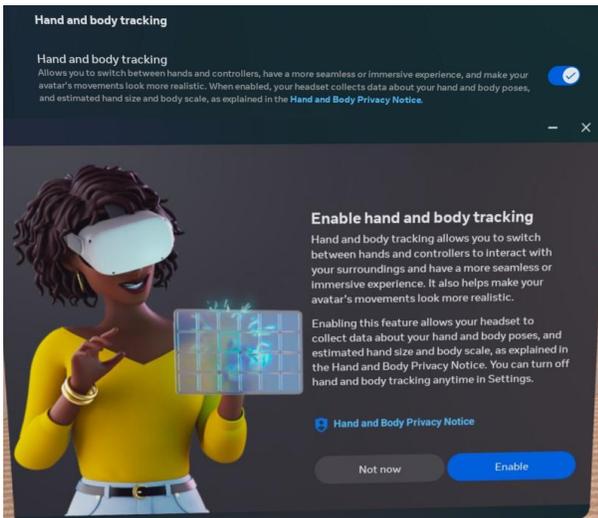


Fig. 11: System-level hand-tracking permission dialog (Oculus)

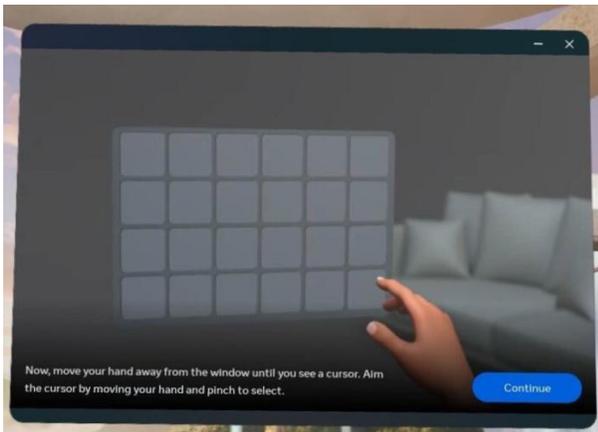


Fig. 12: System-level hand-tracking tutorial dialog (Oculus)

[Q35]: I feel comfortable using the device knowing the level of access it has to my hand-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Instruction] Now, we are interested in the degree to which you understand what the system (i.e., the headset) can do with your data once you grant permission. Answer the following true or false questions regarding the **sensor capability**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q36]: The system can understand your hand gesture to perform certain actions (e.g., select, scroll). 1. True 2. False 3. I don’t know

[Q37]: The system can identify which real-world objects you are holding. 1. True 2. False 3. I don’t know

[Q38]: The system can simulate your hand movement for

your virtual avatar. 1. True 2. False 3. I don’t know

[Q39]: The system can authenticate your identity from the unique aspect of your hand (i.e., fingerprint). 1. True 2. False 3. I don’t know

[Q40]: The system can measure the hand size of new users. 1. True 2. False 3. I don’t know

[Instruction] Answer the following true or false questions regarding the **sensor privacy**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q41]: The system requires your permission to access your hand tracking data. 1. True 2. False 3. I don’t know

[Q42]: The system allows you to control which application has access to your hand tracking. 1. True 2. False 3. I don’t know

[Q43]: The system can transfer your hand tracking data to an external device (e.g., a company server). 1. True 2. False 3. I don’t know

[Q44]: The system can retain the image of your hand on the AR/MR headset. 1. True 2. False 3. I don’t know

[Q45]: The system only collects your final selection (instead of your hand movements) from the hand tracking data. 1. True 2. False 3. I don’t know

[Instruction] Now, you open an app on the headset, which doesn’t need to request app-level permission for hand tracking since the app has automatic access to hand tracking data.

[Q46]: I feel informed about the utility of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q47]: I feel informed about the associated privacy risk of this permission. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q48]: I feel confident that this AR/MR application will securely store my hand-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q49]: I know exactly what type data will be collected, how it will be used, and who will have access to it based on the information presented in the above permission screenshots. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Q50]: I feel comfortable using this AR/MR application knowing the level of access it has to my hand-tracking data. (5pt Likert scale from “Strong disagree” to “Strong agree”)

[Instruction] Now, we are interested in the degree to which you understand what the system (i.e., the headset) can do with your data once you grant permission. Answer the following true or false questions regarding the **sensor capability**. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q51]: The application can understand your hand gesture to perform certain actions (e.g., select, scroll). 1. True 2. False 3. I don't know

[Q52]: The application can identify which real-world objects you are holding. 1. True 2. False 3. I don't know

[Q53]: The application can simulate your hand movement for your virtual avatar. 1. True 2. False 3. I don't know

[Q54]: The application can authenticate your identity from the unique aspect of your hand (i.e., fingerprint). 1. True 2. False 3. I don't know

[Q55]: The application can measure the hand size of new users. 1. True 2. False 3. I don't know

[Instruction] Answer the following true or false questions regarding the sensor privacy. This is not an evaluation of you; rather, we are attempting to evaluate the efficacy of these dialogs.

[Q56]: The application requires your permission to access your hand-tracking data. 1. True 2. False 3. I don't know

[Q57]: The application can access your hand tracking data when running in the background 1. True 2. False 3. I don't know

[Q58]: The application can transfer your hand tracking data to an external device (e.g., a company server). 1. True 2. False 3. I don't know

[Q59]: The application can retain the image of your hand within the application. 1. True 2. False 3. I don't know

[Q60]: The application only collects your final selection (instead of your hand movements) from the hand tracking data. 1. True 2. False 3. I don't know

[Instruction] Now that you have seen the permission settings for both the overall system and the app, we want to understand what information about both the system and the app can help you feel more comfortable using the technology in the future.

[Instruction] Please drag and drop the top three most important items from the list below that can influence your decision to use this technology in the future. (Don't worry about the ordering within the box)

[Item 1]: Knowing who will have access to this data. [Example includes: permission request; background access, control which app has access to your data].

[Item 2]: Knowing how will the data be stored. [Example includes: Delete after use, stores hand tracking data by default; provide options to delete your data.]

[Item 3]: Knowing how will the data be transmitted. [Example includes: keep your data only on device; transfer your data to an external device]

[Item 4]: Knowing what type of data will be collected. [Example includes: hand movement data (how fast you move); hand gesture data (what gesture you perform); unique aspect of your hand (fingerprint).]

[Item 5]: Knowing what is the purpose of collecting this data. [Example includes: indicate selection; generate virtual avatar; identify authentication]

C. Survey for Eye-tracking on HoloLens

[Instruction] Suppose you want to use an AR/MR headset with an eye-tracking feature. Below is what you see in the process of granting permission for eye-tracking. We would like to ask you about your comfort levels and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for eye tracking. You can enable eye-tracking permission, pause eye-tracking, and control eye calibration data for the system in this dialog from the system setting. (Figure 13):

[Instruction] After you enable the eye-tracking feature, you will be asked to perform a calibration process. After the calibration process, the system provides an alternative sign-in process using the eye-tracking feature. This feature is optional (Figure 14):

Questions are identical to Q1-Q15 in Appendix VIII-A

[Instruction] Now, you open an app on the headset, which has its own app-level permission settings for eye tracking. The following app dialog appears after you open the application for the first time (Figure 15):

Questions are identical to Q16-Q30 in Appendix VIII-A

D. Survey for Hand-tracking on HoloLens

[Instruction] Suppose you want to use an AR/MR application with a hand-tracking feature. Below is what you see in the process of granting permission for hand tracking. We would like to ask you about your comfort levels and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for hand tracking. The hand tracking permission for this system is enabled by default. You are informed about the hand-tracking for the system through this visualization. (Figure 16):

[Instruction] Hand tracking does not require calibration from the user. Currently, the system does not offer a way to control

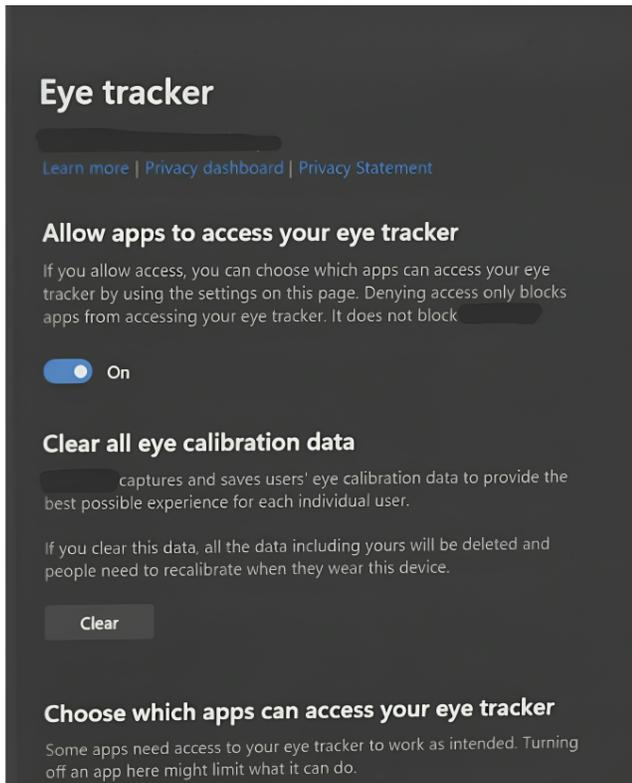


Fig. 13: System-level eye-tracking dialog and app permission control (HoloLens)

which applications can access your hand-tracking data in the system setting.

Questions are identical to Q31-Q45 in Appendix VIII-B

[Instruction] Now, when you open an app on the headset, it does not need to request app-level permission for hand tracking. The hand-tracking permission in this system is enabled by default, so the application automatically has access to the hand-tracking data. You can control the permission for hand-tracking background access for the applications in the system settings (Figure 17):

Questions are identical to Q46-Q60 in Appendix VIII-B

E. Survey for Eye-tracking on Vision Pro

[Instruction] Suppose you want to use an AR/MR application with an eye tracking feature. Below is what you see in the process of granting the permission for eye tracking. We would like to ask you about your comfort level and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for eye tracking. The eye tracking permission for the system is enabled by default. You are informed about the eye-tracking calibration for the system in this dialog. (Figure 18):

[Instruction] After the calibration process, the system provides an alternative sign-in process using the eye-tracking feature.

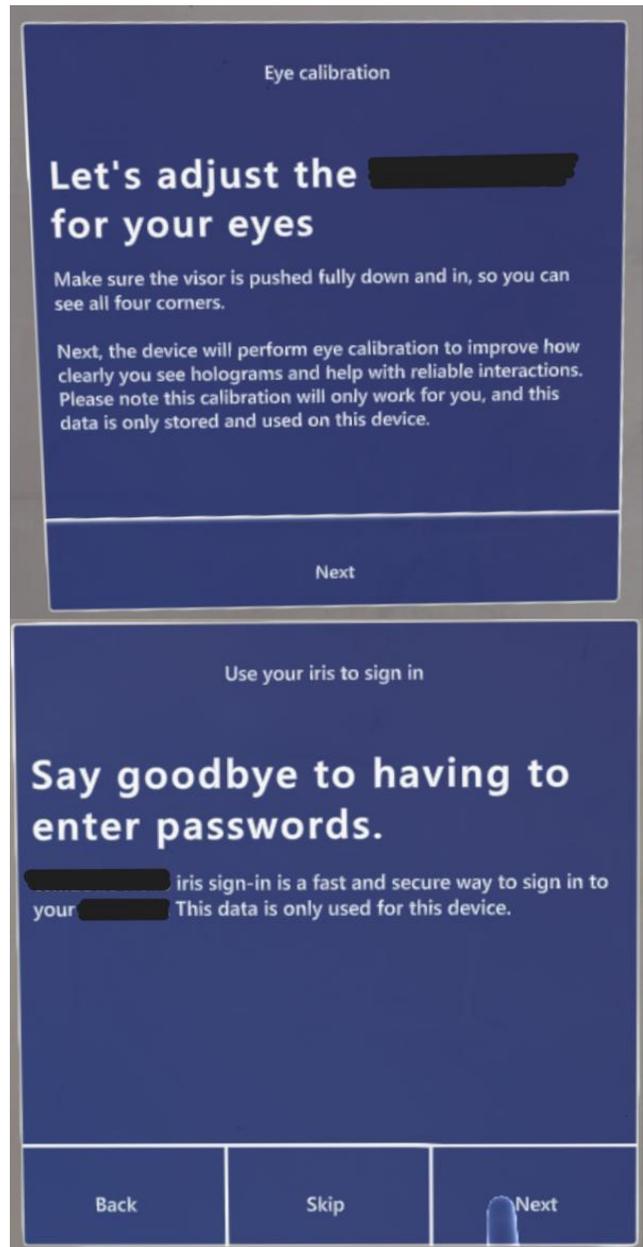


Fig. 14: System-level eye-tracking calibration and iris sign-in dialog (HoloLens)

This feature is optional(Figure 19):

Questions are identical to Q1-Q15 in Appendix VIII-A

[Instruction] Now, you open an app on the headset, which doesn't need to request app-level permission for eye tracking since device doesn't share eye-tracking data with applications.

Questions are identical to Q16-Q30 in Appendix VIII-A

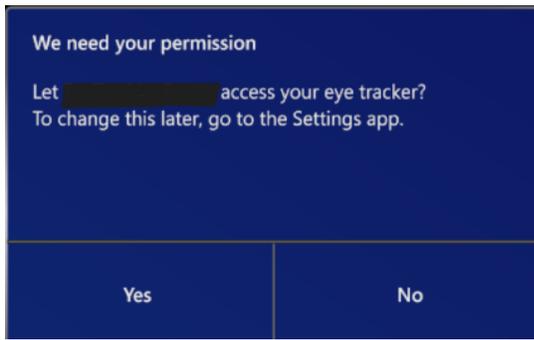


Fig. 15: App-level eye-tracking permission dialog (HoloLens)

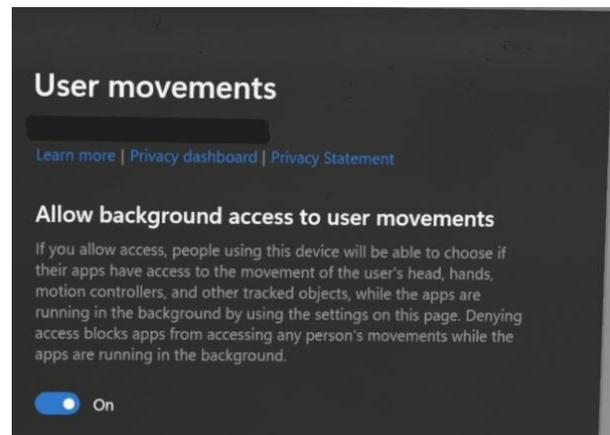


Fig. 17: Background access permission for hand-tracking (HoloLens)

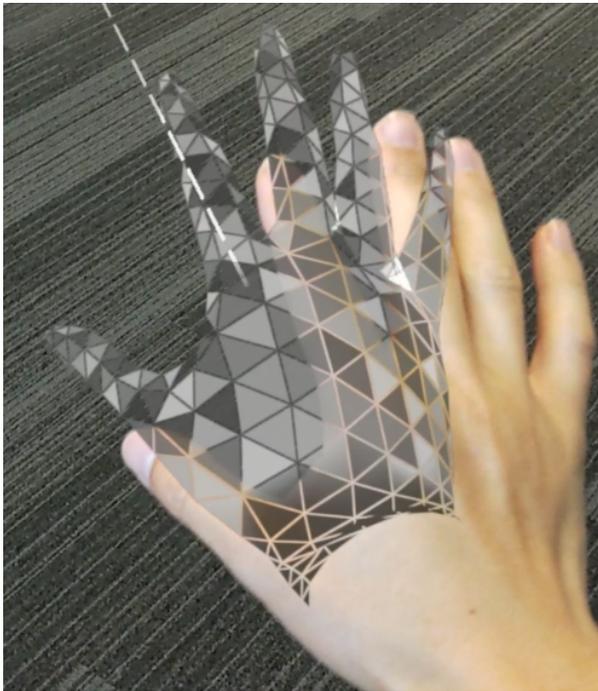


Fig. 16: Hand-tracking visualization (HoloLens)

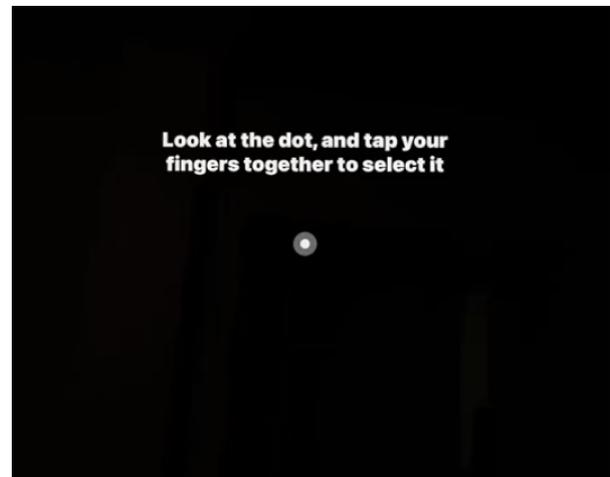


Fig. 18: Eye-tracking calibration (Vision Pro)

F. Survey for Hand-tracking on Vision Pro

[Instruction] Suppose you want to use an AR/MR application with a hand-tracking feature. Below is what you see in the process of granting permission for hand tracking. We would like to ask you about your comfort levels and how informed you feel during this permission-granting flow. You will first navigate the system-level permission settings for hand tracking. The hand tracking permission for this system is enabled by default. You are informed about the hand-tracking for the system through this visualization. (Figure 20):

[Instruction] You can control which application has access to your hand-tracking data in the system setting (Figure 21):

Questions are identical to Q31-Q45 in Appendix VIII-B

[Instruction] Now, you open an app on the headset, which has its own app-level permission settings for hand tracking (Figure 22):

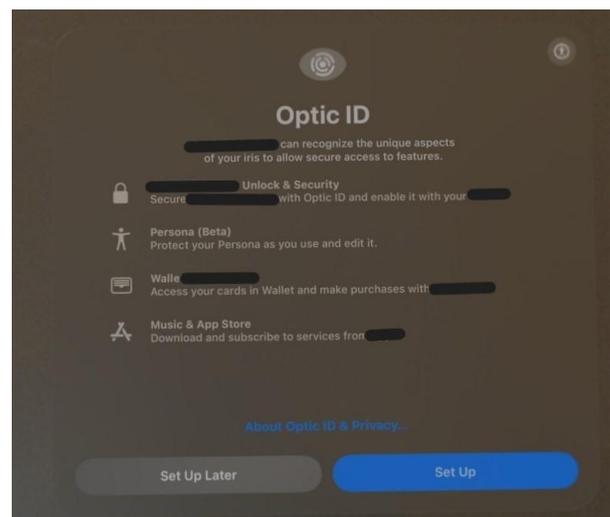


Fig. 19: Eye-tracking Optic ID (Vision Pro)



Fig. 20: Hand-tracking calibration (Vision Pro)

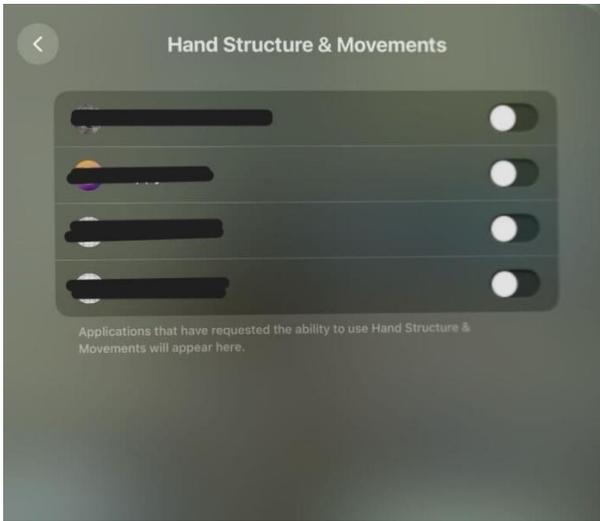


Fig. 21: Hand-tracking app permission control (Vision Pro)

Questions are identical to Q46-Q60 in Appendix VIII-B

IX. COMPREHENSION QUESTIONS ANSWER KEY

As part of our analysis of participant comprehension, we determine our own best assessment of the correct answer. We did this based on our own understanding of the APIs, documentation, and privacy policies. We document our answers and justifications for eye-tracking on the system level in Table V, eye-tracking on the application level in Table VI, hand-tracking on the system level in Table VII, and hand-tracking on the application level in Table VIII. Quotes from privacy policies or documentation are in italics in the tables.

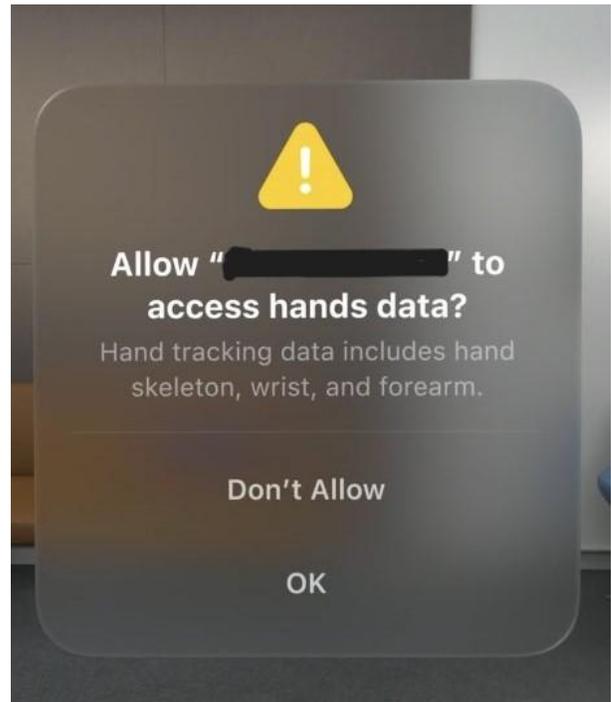


Fig. 22: App-level hand-tracking permission dialog (Vision Pro)

TABLE V: Justification for system-level eye-tracking comprehension questions

Comprehension Question	HoloLens	HoloLens Justification	Oculus	Oculus Justification	Vision Pro	Vision Pro Justification
The system requires your permission to access your eye-tracking data.	False	Denying access only blocks apps from accessing your eye tracking. It does not block HoloLens (Figure 13)	True	Run-time system level permission model (Figure 8)	False	Eye-tracking is enabled by default for the system.
The system allows you to control which application has access to your eye tracking.	True	Built-in function (Figure 13)	True	Built-in function (Figure 9)	False	Eye input is not shared with Apple, third-party apps, or websites. [31]
The system can transfer your eye-tracking data to an external device (e.g., a company server).	False	Microsoft doesn't store any biometric or other identifiable information [29].	True	We collect and retain certain data about your interactions with eye tracking [14]	False	Eye input is not shared with Apple, third-party apps, or websites. [31]
The system can retain the unprocessed image of your eye on the AR/MR headset.	False	We store calibration information locally on device correlated with bit codes from the Iris pattern [26]	False	The raw image data is deleted from your headset after the abstracted gaze data is generated. [14]	False	Optic ID data — including mathematical representations of your iris — is encrypted and protected by the Secure Enclave [31]
The system only collects your final selection (instead of your eye movements) from the eye tracking data.	False	Abstracted eye-tracking data is available to the system [20]	False	Abstracted eye-tracking data is available to the system [10]	True	Data minimization for eye-tracking data [31]
The system can understand where your eyes look to indicate which virtual object to select.	True	Built-in function (EyesPose.Gaze [20]).	True	Built-in function (OVREyeGaze [10]).	True	Built-in function [31].
The system can identify which real-world objects you are looking at.	True	Access to passthrough camera data is available [5].	False	Identifying real-world objects requires integrating passthrough camera data, which the eye-tracking API does not offer [2].	False	Identifying real-world objects requires integrating passthrough camera data, which the eye-tracking API does not offer [31].
The system can simulate your eye movement for your virtual avatar.	True	Abstracted eye-tracking can simulate eye movement (EyesPose.Gaze [20]).	True	Abstracted eye-tracking can simulate eye movement (OVREyeGaze [10]).	True	Built-in function (Persona [31]).
The system can authenticate your identity from the unique aspect of your eye (i.e., iris).	True	Store calibration information locally on device correlated with bit codes from the Iris pattern [26]	False	Iris-scanning function is not supported	True	Optic ID data is encrypted, never leaves your device, and is accessible only to the Secure Enclave processor. [31]
The system can adjust eye calibration for new users.	True	Built-in function [26].	True	Built-in function [10]).	True	Built-in function [31].

TABLE VI: Justification for application level eye-tracking comprehension questions

Comprehension Question	HoloLens	HoloLens Justification	Oculus	Oculus Justification	Vision Pro	Vision Pro Justification
The application requires your permission to access your eye-tracking data.	True	Run-time app level permission model (Figure 15)	True	Run-time app level permission model (Figure 10)	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application can access your eye-tracking data when running in the background.	False	Background access for eye-tracking is not supported.	False	Background access for eye-tracking is not supported.	False	Eye input is not shared with Apple, third-party apps, or websites. [31]
The application can transfer your eye-tracking data to an external device (e.g., a company server).	True	System provide no control over how third-party used user’s eye-tracking data .	True	Oculus does not control how a third-party app uses, stores, or shares your abstracted gaze data [14]	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application can retain the unprocessed image of your eye on the AR/MR headset.	False	Only abstracted eye-tracking data is available to the application [20]	False	Only abstracted eye-tracking data is available to the application [14]	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application only collects your final selection (instead of your eye movements) from the eye-tracking data.	False	Application has access to the abstracted eye-tracking data [20]	False	Application has access to the abstracted eye-tracking data [10]	True	Only when you select the button, by both looking at it and tapping your fingers together, does where you are looking get communicated to the app. [31].
The application can understand where your eyes look to indicate which virtual object to select.	True	Built-in function (EyePose.Gaze [20]).	True	Built-in function (OVREyeGaze [10]).	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application can identify which real-world objects you are looking at.	True	Access to passthrough camera data is available [5].	False	Identifying real-world objects requires integrating passthrough camera data, which the eye-tracking API does not offer [2].	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application can simulate your eye movement for your virtual avatar.	True	Abstracted eye-tracking can simulate eye movement (EyePose.Gaze [20]).	True	Abstracted eye-tracking can simulate eye movement (OVREyeGaze [10]).	False	Eye input is not shared with Apple, third-party apps, or websites [31].
The application can authenticate your identity from the unique aspect of your eye (i.e., iris).	False	All calibration data is stored securely on the device locally and only available to the system [26]	False	Iris-scanning function is not supported	False	Optic ID data is encrypted, never leaves your device, and is accessible only to the Secure Enclave processor. [8]
The application can access user’s eye calibration data (e.g., eye position) provided by the system.	False	All calibration data is stored securely on the device locally and only available to the system. [26]	False	The eye-tracking API may only request eye-tracker calibration instead of directly accessing the data [34].	False	Data used to calibrate your Apple Vision Pro to your eyes is protected on-device [31].

TABLE VII: Justification for system-level hand-tracking comprehension questions

Comprehension Question	HoloLens	HoloLens Justification	Oculus	Oculus Justification	Vision Pro	Vision Pro Justification
The system requires your permission to access your hand-tracking data.	False	Hand-tracking is enabled by default for the system.	True	Run-time system level permission model (Figure 11)	False	Hand-tracking is enabled by default for the system.
The system allows you to control which application has access to your hand tracking.	False	System automatically grants applications access to the hand-tracking API	False	System automatically grants applications access to the hand-tracking API	True	Run-time application level permission model (Figure 21)
The system can transfer your hand-tracking data to an external device (e.g., a company server).	False	HoloLens also detects hand gestures intended for system interactions (such as menu navigation, pan/zoom, and scroll). This data is processed on your HoloLens device and is not stored. [33].	True	Meta processes and shares the hand-tracking data with the Oculus server, where it is retained for 90 days [15]	False	Apps do not need access to your hands set up information in order to help you interact with content [31].
The system can retain the unprocessed image of your hand on the AR/MR headset.	False	HoloLens also detects hand gestures intended for system interactions (such as menu navigation, pan/zoom, and scroll). This data is processed on your HoloLens device and is not stored. [33].	False	All of this analysis is done on your device in real-time as you move, and the images and estimated points are deleted in real time after processing. We do not collect or store this data on Meta servers [15].	False	Apple Vision Pro measures and stores information on-device about the size and shape of your hands and finger joints to make it easier for you to interact with content [31].
The system only collects your final selection (instead of your hand movements) from the hand tracking data.	False	Abstracted hand-tracking data is available to the system [4]	False	Abstracted hand-tracking data is available to the system [11]	False	Abstracted hand-tracking data is available to the system [19]
The system can understand your hand gestures to perform certain actions (e.g., select, scroll).	True	Built-in function [3]).	True	Built-in function [16].	True	Built-in function [17].
The system can identify which real-world objects you are holding.	True	Access to passthrough camera data is available [5].	False	Identifying real-world objects requires integrating passthrough camera data, which the hand-tracking API does not offer [11].	False	Identifying real-world objects requires integrating passthrough camera data, which the hand-tracking API does not offer [31].
The system can simulate your hand movement for your virtual avatar.	True	Built-in function [4]).	True	Built-in function [11].	True	Built-in function (Persona [31]).
The system can authenticate your identity from the unique aspect of your hand (i.e., fingerprint).	False	Fingerprint authentication function is not supported	False	Fingerprint authentication function is not supported	False	Fingerprint authentication function is not supported
The system can measure the hand size of new users.	True	Built-in function [4]).	True	Built-in function [11].	True	Built-in function [19].

TABLE VIII: Justification for app-level hand-tracking comprehension questions

Comprehension Question	HoloLens	HoloLens Justification	Oculus	Oculus Justification	Vision Pro	Vision Pro Justification
The application requires your permission to access your hand-tracking data.	False	System automatically grants applications access to the hand-tracking API	False	System automatically grants applications access to the hand-tracking API	True	Run-time app level permission model (Figure 22)
The application can access your hand-tracking data when running in the background.	True	Built-in function (Figure 17)	False	Background access for hand-tracking is not supported.	False	Background access for hand-tracking is not supported.
The application can transfer your hand-tracking data to an external device (e.g., a company server).	True	System provides no control over how third-party used user's eye-tracking data	True	...we do not control how a third party app uses, stores, or shares your abstracted hand and body data. [15]	True	It's [developer] responsibility to protect any data your app collects, and to use it in responsible and privacy-preserving ways [7]
The application can retain the unprocessed image of your hand on the AR/MR headset.	True	Access to passthrough camera data is available [5].	False	Only abstracted hand-tracking data is available to the application [11].	False	Only abstracted hand-tracking data is available to the application [19].
The application only collects your final selection (instead of your hand movements) from the hand-tracking data.	False	Abstracted hand-tracking data is available to the application [4]	False	Abstracted hand-tracking data is available to the application [11]	False	Abstracted hand-tracking data is available to the application [19]
The application can understand your hand gestures to perform certain actions (e.g., select, scroll).	True	Built-in function [3]).	True	Built-in function [16].	True	Built-in function [17].
The application can identify which real-world objects you are holding.	True	Access to passthrough camera data is available [5].	False	Identifying real-world objects requires integrating passthrough camera data, which the hand-tracking API does not offer [11].	False	Identifying real-world objects requires integrating passthrough camera data, which the hand-tracking API does not offer [31].
The application can simulate your hand movement for your virtual avatar.	True	Built-in function [4]).	True	Built-in function [11].	True	Built-in function [19].
The application can authenticate your identity from the unique aspect of your hand (i.e., fingerprint).	False	Fingerprint authentication function is not supported	False	Fingerprint authentication function is not supported	False	Fingerprint authentication function is not supported
The application can analyze the hand size of new users.	True	Built-in function [4]).	True	Built-in function [11].	True	Built-in function [19].