







# Visceral Notices and Privacy Mechanisms for Eye Tracking in Augmented Reality

Nissi Otoo\* , Kailon Blue\* , G. Nikki Ramirez , Evan Selinger , Shaun Foster , and Brendan David-John 

**Abstract**—Head-worn augmented reality (AR) continues to evolve through critical advancements in power optimizations, AI capabilities, and naturalistic user interactions. Eye-tracking sensors play a key role in these advancements. At the same time, eye-tracking data is not well understood by users and can reveal sensitive information. Our work contributes visualizations based on visceral notice to increase privacy awareness of eye-tracking data in AR. We also evaluated user perceptions towards privacy noise mechanisms applied to gaze data visualized through these visceral interfaces. While privacy mechanisms have been evaluated against privacy attacks, we are the first to evaluate them subjectively and understand their influence on data-sharing attitudes. Despite our participants being highly concerned with eye-tracking privacy risks, we found 47% of our participants still felt comfortable sharing raw data. When applying privacy noise, 70% to 76% felt comfortable sharing their gaze data for the Weighted Smoothing and Gaussian Noise privacy mechanisms, respectively. This implies that participants are still willing to share raw gaze data even though overall data-sharing sentiments decreased after experiencing the visceral interfaces and privacy mechanisms. Our work implies that increased access and understanding of privacy mechanisms are critical for gaze-based AR applications; further research is needed to develop visualizations and experiences that relay additional information about how raw gaze data can be used for sensitive inferences, such as age, gender, and ethnicity. We intend to open-source our codebase to provide AR developers and platforms with the ability to better inform users about privacy concerns and provide access to privacy mechanisms. A pre-print of this paper and all supplemental materials are available at [https://bmdj-vt.github.io/project\\_pages/privacy\\_notice](https://bmdj-vt.github.io/project_pages/privacy_notice).

**Index Terms**—Augmented reality, Eye tracking, Privacy notice

## 1 INTRODUCTION

Eye tracking enables exciting developments in augmented reality (AR), making experiences more interactive and intuitive by leveraging gaze behavior. However, eye tracking raises privacy concerns based on how personal gaze data is collected, used, and potentially exploited, whether or not users have control over this sensitive information [30]. Eye-tracking data continues to become more ubiquitous in extended reality (XR) research and applications [11, 11, 56], suggesting a future where inferences from gaze data are enabled in exchange for enhanced capabilities in optimized rendering [22, 50], natural interactions [42], or AI assistance [10].

Typical AR users lack a robust or coherent understanding of their eye-tracking data compared to developers and XR enthusiasts [1, 2]. Existing work has explored Visceral Interfaces (VIs) to describe an experientially rich approach for relaying privacy-relevant information through sensory and psychological cues [12]. Visceral notices were conceptualized to enhance user awareness within notice and consent and leverage user familiarity with experiences or symbolism, expected psychological responses to stimuli, and by showing implications or results of their decisions. These dimensions provide novel approaches to leverage visualizations and experiences in immersive settings that relay privacy implications. VIs specific to privacy awareness for Virtual Reality (VR) eye tracking have been proposed [48] and evaluated [43], demonstrating their value for increasing privacy awareness for this emerging technology.

Eye-tracking data can reveal not only where someone is looking, but also sensitive information such as personality traits, sexual preferences,

biases, substance use, and medical conditions [32]. This can lead to highly invasive profiling [20] that is highly unsettling and disturbing to users [51]. To address these concerns, privacy mechanisms have been developed to degrade eye-tracking data [9, 20, 51] and reduce the risk of sensitive inferences from untrusted parties. At the same time, privacy notices and visualizations are capable of informing users of what their gaze data reveals about them. While visualizations make data more comprehensible to users, it is increasingly difficult for users to understand the role of highly technical privacy mechanisms and their ability to mitigate the risks of sharing eye-tracking information [18, 19, 51, 53]. No research on privacy mechanisms or VIs has been conducted in an AR environment, nor have visualizations been used to inform users about the effect of gaze-based privacy-preserving mechanisms on their data. Our work unifies these efforts by leveraging VIs to understand subjective perceptions of the safety provided by privacy mechanisms applied in head-worn AR settings.

Our work presents two main contributions. First, we adapt gaze-based VIs from VR to AR and extend their evaluation beyond free viewing to a gaze-based selection task. The findings support existing work from VR on user preferences towards specific forms of visualizations while providing new insights when gaze data is actively applied to the user's current task. With the jump from VR to AR, we also consider that VIs provide insight into user gaze behavior with both real and virtual content, instead of only virtual content the device is already aware of. Second, we are the first to leverage VIs to visualize the effect of privacy mechanisms on gaze data streams and understand subjective attitudes towards these techniques and their impact on data-sharing preferences. These insights complement existing quantitative analysis on how well these mechanisms preserve privacy and further support the use of Weighted Smoothing to ensure users are comfortable sharing data while they are protected against privacy attacks such as re-identification.

## 2 RELATED WORKS

### 2.1 Eye-Tracking Privacy

Eye-tracking data enables practical applications within XR systems but also presents the risk of leaking personal and sensitive information when shared with others [11]. Gaze data serves as a biometric capable of identifying a user with high accuracy in both desktop [39] and XR settings [5, 38], posing a risk for re-identification [20]. Fea-

- \*Both authors contributed equally.
- Nissi Otoo is with Virginia Tech. E-mail: [nissiotoo@vt.edu](mailto:nissiotoo@vt.edu).
- Kailon Blue is with Virginia Tech. E-mail: [kailonb@vt.edu](mailto:kailonb@vt.edu).
- G. Nikki Ramirez is with Virginia Tech. E-mail: [gnram@vt.edu](mailto:gnram@vt.edu).
- Evan Selinger is with Rochester Institute of Technology.
- Shaun Foster is with Rochester Institute of Technology.
- Brendan David-John is with Virginia Tech. E-mail: [bmdj@vt.edu](mailto:bmdj@vt.edu).

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: [reprints@ieee.org](mailto:reprints@ieee.org).  
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

tures extracted from gaze data are also capable of profiling physical and behavioral traits such as age [41], race [7], gender [46], and ad preferences [28]. Eye-tracking data poses a unique risk in that reactive eye movements cannot be controlled; even informed users are unable to consciously control their gaze or limit what it reveals about them [32].

Privacy-enhancing technologies (PETs) specific to gaze data have been developed to protect user data by degrading quality and fidelity as a trade-off between utility and privacy. Several approaches are applied directly to streams of gaze data [53], including access control for safeguarding data with a gatekeeper only releasing the data required for specific applications [20] and content-aware methods that provide a spatial differential privacy (DP) guarantee [36]. Privacy protections on data streams are limited in that practical adoption of the approach is unknown and not legally required, resulting in applications that primarily rely on notice and consent to enable and track gaze data. Other approaches support privacy-preserving dataset sharing through formal privacy guarantees, either through differential privacy [9, 51] or alternative guarantees specific to re-identification [18, 19]. The privacy-utility trade-off resulting from these PETs is primarily evaluated with offline data analysis against privacy attacks and does not measure user perceptions or attitudes towards their data after PETs are applied. Steil et al. measured data-sharing attitudes but did not measure the impact on attitudes after their DP privacy mechanism was applied [51]. Li et al. explored user preferences through interactions with slider controlling privacy noise level within an interactive game, but did not link DP parameters to data-sharing attitudes or privacy awareness [36]. Our work provides a new understanding of user preferences in gaze-based privacy mechanisms and informs privacy-utility trade-offs in an AR context.

## 2.2 Privacy and Security Indicators

Privacy policies are notoriously difficult for users to understand and are often filled with dense language and legal jargon, primarily protecting companies from liability rather than genuinely informing users. As a result, most people either skim these policies or avoid reading them altogether, leaving them unaware of the full scope of data being collected and how it might be used [31, 47, 52]. Users prefer simplicity in privacy notices, favoring plain language and transparency over verbose and legalistic descriptions [44]. This lack of clarity in data collection creates a disconnect between users and how their data is processed, motivating the need for effective privacy and security indicators to support large-scale adoption of XR systems and the invasive sensors they include [21].

Privacy and security indicators are in development for VR [37, 43, 55] but lack evaluation in an AR context. AR data collection is increasingly pervasive and personal [26] and should offer real-time, visually engaging notifications to inform users. For example, researchers have suggested using visual elements such as color-coded outlines around avatars in social settings to confirm their authenticity [37] or icons above virtual portals to indicate connections to other parts of the Metaverse [55]. By integrating these real-time, intuitive notifications, AR systems can raise awareness of data collection while maintaining a positive user experience. The current standard in eye-tracking privacy notice for AR comes from the Magic Leap 2 data transparency policy [35] which includes the requirement that “Users must be able to visualize (e.g., a red dot recording logo) when eye tracking data is being collected, stored, transferred or otherwise used by the Application.” It remains an open question as to how many apps currently follow this guidance and what visualization options they provide. Our work provides access to active AR visual indicators beyond adaptations of 2D indicators.

## 2.3 Visceral Notices for Eye-Tracking Data

Visceral notices provide a method of informing users of privacy risks through stimulating, emotionally resonant, and vivid methods [12]. Visceral notices are intended to be intuitive and informative by leveraging stimuli users are familiar with to elicit psychological responses. Rumble strips on the expressway are a real-world example: immediate haptic feedback is used to quickly inform driver safety. Another

example is adding a pair of eyeballs to a tip jar to elicit the feeling of being watched and increase tip amounts. Calo proposed the use of visceral notice in the context of notice and consent. By allowing users to experience information that is typically revealed in lengthy data policies, visceral notices provide attention-grabbing and digestible information that may benefit user autonomy. Selinger et al. proposed multiple VIs to inform VR users when eye tracking is taking place and what data is being captured [48]. For example, being watched by a virtual character as you interact with a VR environment could induce the feeling of being watched and relay a reminder that your movements are being captured. Other proposed VIs explore the concept of showing: data is presented after viewing statues in an art gallery, the proportion of gaze spent on various skin tones is identified, and the user is informed of what their behavior reveals about them.

Ramirez-Saffy et al. were the first to evaluate VIs in a VR setting to enhance privacy awareness around eye tracking [43]. They deployed two VIs. First, a tendrill approach that draws a line indicating past gaze positions behind a crosshair, illustrating the user’s gaze in real-time as it moves around the VR environment. Second, they evaluated an eye icon interface mirroring tracked data as a floating pair of eyeballs in the user’s periphery. We provide more details and an illustration of these visualizations in Section 3.2. The authors studied perceptions of these visualizations in a free-viewing VR art gallery. Their results indicated that the tendrill interface, although somewhat distracting, was more effective at helping users understand gaze behavior and become more aware of how it interacted with their surroundings. The eye icon, while less intrusive, provided fewer insights into the user’s eye movements. The results revealed clear strengths and weaknesses for each interface and significantly lowered willingness to share data with certain entities after experiencing the VIs. Our study extends the evaluation of VIs to AR and includes an active gaze task to measure the trade-off between these VIs and the impact of privacy mechanisms.

**Summary** While previous research has made progress in understanding eye-tracking privacy risks [32], developing privacy mechanisms [9, 20, 51, 53], and exploring visceral interfaces in VR settings [43, 48], there are still key gaps in how these approaches apply to AR environments. First, the impact of visibility from see-through AR displays on visualization preferences and attitudes is not known, and the impact of seeing gaze data interact with real-world content beyond simulated VR environments on data-sharing attitudes is not established. Second, current AR data-sharing policies, such as the Magic Leap 2 [35], require visual indicators for eye tracking; no studies have evaluated the effectiveness of visualizations to balance user experience and privacy awareness for AR eye-tracking data. Similarly, while privacy mechanisms have been tested against technical attacks [18, 19, 53], there is a lack of understanding of how users perceive these mechanisms or how they affect data-sharing attitudes once visualized. For example, Steil et al. [51] researched data-sharing attitudes via surveys but did not measure how they changed after applying their differential privacy mechanism. Li et al. [36] explored privacy noise controls for user experience but did not measure privacy awareness or attitudes.

## 3 METHODOLOGY

We conducted a user study to answer the following questions:

- **RQ1:** What impact does experiencing gaze-based VIs and privacy mechanisms have on privacy attitudes and data-sharing sentiments for AR gaze data?
- **RQ2:** Are user preferences towards AR tendrill and icon VIs and their deployment consistent with VR findings?
- **RQ3:** How do users perceive visualized privacy mechanisms and what are their preferences towards sharing gaze data with or without privacy mechanisms applied?

Pre and post-experiment surveys measured the impact of experiencing gaze-based VIs and privacy mechanisms on data-sharing attitudes and sentiments across two AR tasks (**RQ1**). Surveys on user preferences for VI type and types of deployments in practice evaluated whether existing results from VR settings are consistent in AR (**RQ2**) along with user experience metrics and subjective preferences toward three privacy mechanisms for AR gaze data (**RQ3**).

### 3.1 Protocol

Our user study was approved by our Institutional Review Board (IRB) and subjects were paid \$20 USD. Studies took 40 minutes on average for the Art Gallery and 55 minutes on average for the Gaze Selection task. Participants were recruited through email and word of mouth. The user study flow is outlined in Figure 1. Exact survey questions and measures are provided in the Supplementary Material for reproducibility.

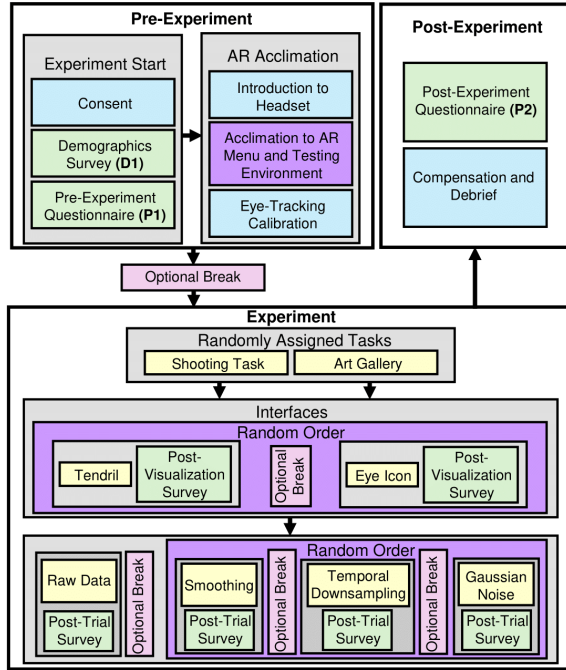


Fig. 1: User study protocol. Participants were split into groups based on task (Gaze Selection and Art Gallery) with each visualization interface in a counter-balanced order. Participants experienced raw data visualized as a baseline then all three privacy mechanisms in a counter-balanced order for each interface.

#### 3.1.1 Demographics

The study consisted of 34 participants with no specific technical expertise required. We collected basic demographic information, including gender identity, race/ethnicity, age, and prior experience with XR devices. Of the participants, 14 identified as Caucasian or White, 12 as Asian, 4 as Black or African American, 2 as Multiracial, 1 as Moroccan, and 1 as Hispanic or Latino. Regarding gender, 24 identified as men, 9 as women, and 1 as non-binary. Participants were between the ages of 18 and 30 with a mean age of 20.8. Prior XR experience and levels of privacy concern were also measured between participants (Table 1). Most participants (82.4%) had prior VR experience with 82.1% of these participants reporting less than three hours of VR activity in the past month. For privacy concerns, 10 reported their medical information had been improperly disclosed, with 3 indicating they were victims of an improper invasion of privacy. A total of 6 participants indicated they were victims of an improper invasion of privacy while 8 others stated they did not know if they were. Most (26) participants agreed consumers do not have control over how personal information is collected, 18 of those participants did not identify as victims; 5 participants agreed most businesses handle such information properly and confidentially. Only 4 participants did not agree with either statement and were either unsure or not victims of privacy invasion.

#### 3.1.2 Pre-Experiment

After providing demographic data, participants viewed a short presentation from the researcher on the role of privacy mechanisms in balancing privacy and utility, definitions of our privacy mechanisms, and a video recording showing raw data and data after each privacy mechanism was

applied. The objective of the presentation was to provide a uniform understanding for participants on privacy mechanism implementation and did not include any information comparing how effective they were in protecting privacy or their impact on utility. To avoid biasing participants, the presentation did not discuss how gaze data creates privacy risks as in related work [43] and did not emphasize any of the three possible risks (re-identification, targeted ads, and sensitive inferences) that were listed as examples of risks privacy mechanisms are designed to mitigate while balancing utility in general. A pre-experiment survey then measured participants' initial feelings about eye-tracking data sharing and how comfortable they were with sharing their gaze data for different purposes and with different entities. These responses provided a baseline understanding of their privacy attitudes prior to experiencing VIs to answer RQ1.

#### 3.1.3 Experiment

Once the surveys were completed, participants were randomly assigned to one of two task groups: Art Gallery and Gaze Selection.

**Art Gallery:** Participants freely explored an AR art gallery aligned with the walls of a long hallway (Figure 2, Left). They were tasked with using a hand controller to select their five favorite art pieces within a three-minute time limit. This task replicates an environment where eye-tracking data is collected passively to monitor user behavior and mirrors an existing VI evaluation task [43].




Fig. 2: Left: Art Gallery task. Right: Gaze Selection task with Eye Icon visualization in the top-right of the field of view.

**Gaze Selection:** Participants were presented with a gaze-based target selection task in the presence of distractors (Figure 2, Right). Cartoon animals—squirrels or cats—randomly spawn in front of the user for one minute, gaining points for selecting the squirrels and losing points when selecting cats. Dwell-time was used to select targets with a selection triggered by maintaining gaze for 500 ms. This task replicates environments where users' eye-tracking data is actively used within the AR application and was motivated by the shooting task from prior privacy mechanism evaluations in VR [53]. In privacy mechanism conditions, the privacy-enhanced gaze data was used for selection.

Data was collected across the two tasks in a between-subjects fashion with 18 participants for Art Gallery and 16 for Gaze Selection. Participants experienced their assigned task using both VIs presented in a counter-balanced order. Within each VI block, participants experienced trials with visualizations of raw gaze data followed by the three privacy mechanisms counter-balanced.

#### 3.1.4 Post-Block

After each VI block, participants completed a survey to measure their workload using the NASA-TLX questionnaire [27] and attitudes towards the visualization technique. This survey also collected their attitudes and comfort levels toward sharing their gaze data using the same questions from the pre-experiment survey. Additionally, participants rated how comfortable they were in sharing data across the privacy mechanisms. These responses provide an understanding of VI along several dimensions for answering RQ2. We also measured motion sickness using a questionnaire from the National Library of Medicine [8, 13].



Table 1: Demographic survey questions related to prior XR experience and levels of privacy concern.

Demographic	Response Format
Have you used a virtual-reality headset before?	Yes/No
Estimate the number of hours you have used head-mounted virtual or augmented reality in the last month (e.g., Oculus/Meta Quest).	Open-ended
Estimate the number of hours you have used 3D applications (including video games) in the last month.	Open-ended
Which of the following applications have you used virtual reality for?	(1) Social VR Learning/Education (2) Gaming (3) Cinematic Experiences (4) Streaming Live Sports/ Entertainment (5) Have not used virtual reality before
Privacy Concern Questions	Response Format
Which of the following do you believe has ever disclosed your personal medical information in a way that felt improper?	(1) Health insurance companies (2) A clinic or hospital that treated you or a family member (3) Public health agencies (4) Your employer or a family member's employer (5) A doctor who has treated you or a family member (6) A pharmacist who filled a prescription for you or a family member
Have you personally ever been the victim of what you felt was an improper invasion of privacy, or not?	(1) Yes, I have been a victim (2) No, I have not been a victim (3) I don't know
Which of these statements do you think are true?	(1) Consumers have lost all control over how personal information is collected and used by companies. (2) Most businesses handle the personal information they collect about consumers in a proper and confidential way. (3) N/A

### 3.1.5 Post-Trial

Participants completed a survey after each trial regarding the visualization of raw gaze data and the three privacy mechanisms. This survey collected their feelings regarding their ability to complete their task, perceived privacy with the mechanism (if applicable), and willingness to share the data collected when the privacy mechanism was applied for each condition. The post-trial surveys were based on the User Experience Questionnaire (UEQ) as a standard scale for evaluating interfaces [34]. We omitted certain UEQ questions that did not fit the scope of this study based on feedback from pilot testing. The final set of questions were aggregated to compute scores for attractiveness, perspicuity, efficiency, dependability, stimulation, and novelty. These responses provide a comparison of preferences and perceptions of the explored privacy mechanisms for aggregate and between-task analysis to answer **RQ3**.

### 3.1.6 Post-Experiment

Participants completed a post-experiment survey mirroring the pre-experiment survey. The goal of this survey was to assess whether their views on sharing their eye-tracking data had changed holistically after interacting with both VIs and the privacy mechanisms.

## 3.2 Visualizations

The interfaces were implemented on a Magic Leap 2 headset using Unity 2022.3.30f1 and adapted from the open-source code provided by Ramirez-Saffy et al. [43]. Interface animations were based on real-time gaze data from the Magic Leap 2 sampled at 60 Hz.

**Tendrill Interface:** The current gaze position in 3D space was used to draw a crosshair with a trail renderer component that follows gaze over time. The tendrill interface constantly displays the current gaze embedded within the real world. However, by appearing wherever a user looks, the tendrill interface is often considered distracting and annoying despite being preferred to inform privacy awareness in VR [43]. The tendrill VI is illustrated for raw and privacy-enhanced gaze data in Figure 3.

**Eye Icon Interface:** The eye icon interface offers a more abstract approach to visualizing gaze data by animating a set of virtual eyes located at a fixed position in the upper-right quadrant of the user's field of view. This interface uses a skeuomorphic design of animated eyes that makes use of the "familiarity as warning" dimension of visceral notice [48]. The goal of the VI is to provide a persistent reminder that their gaze activity is being recorded without distracting them. The interface takes advantage of familiarity with recording indicators such as a blinking red circle but was adapted to rotate with the user's eye movements as they occur.

We modified existing implementations [43] in two ways. First, participants in their study reported that the main limitation of the eye

icon is that the user can never actually see the eye movements: because gaze data was sampled quickly, by the time they glanced at the icon, users only observed static gaze in the direction of the icon position. We introduced a 250ms delay in animation from the gaze data stream to allow users to observe a short window into past eye movements. Second, participants reported that the eye icon was creepy and overwhelming. While this captures the goal of the VI, we wanted to explore if adjusting the icon based on context would reduce the severity of this effect. The floating eyeballs were maintained in the Art Gallery task while the animated eyes were applied to a floating cat in the Gaze Selection task to align with their selection targets (Figure 2, Right).

## 3.3 Privacy Mechanisms

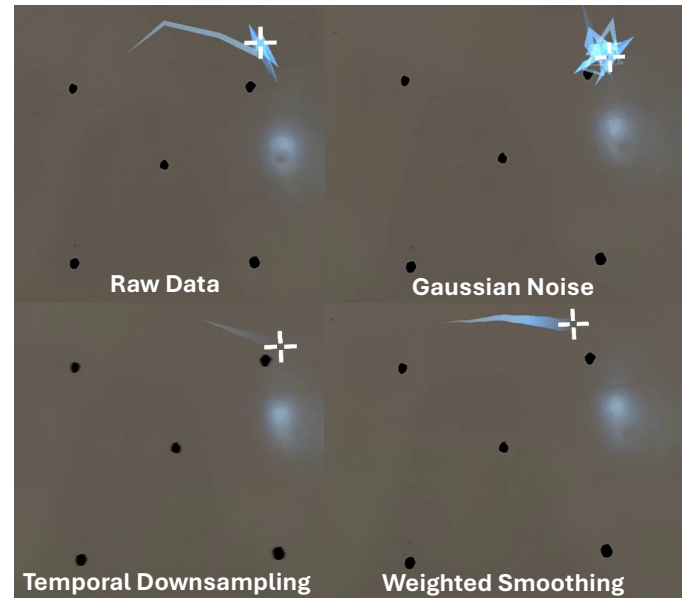


Fig. 3: Raw gaze data and privacy mechanisms using the tendrill interface as gaze shifts toward the top-right target.

State-of-the-art analysis of privacy-utility trade-offs for streamed gaze data has identified Gaussian Noise ( $\sigma = 1^\circ$  and  $\sigma = 3^\circ$  as low and high Gaussian Noise, respectively), Temporal Downsampling, and Weighted Smoothing as the top-performing mechanisms for reducing the risk of re-identification [53]. Figure 3 illustrates the three gaze privacy mechanisms with the tendrill interface to demonstrate their influence on gaze data. Participants were randomly assigned privacy mechanisms in a counter-balanced order after experiencing the raw

data trail in each VI block. Privacy mechanisms are applied to the gaze samples  $X_n$  to produce a perturbed sample  $X'_n$  where  $n$  is the index of the gaze sample. Gaze is represented within the eye-in-head coordinate frame using spherical coordinates [23] that represent horizontal and vertical gaze angles with a unit radius of one  $X_n = (\theta_n, \psi_n, r_n = 1)$ . The privacy mechanisms were implemented as follows:

**Gaussian Noise:** Previous research on low and high rates reduces the re-identification rates to 40% ( $\sigma = 1^\circ$ ) and 20% ( $\sigma = 3^\circ$ ) [53]. Coordinates  $\theta_n, \psi_n$  of  $X_n$  are offset by noise randomly sampled from a normal distribution with zero mean and a standard deviation of  $\sigma = 1.5^\circ$  to retain utility for target selection and ensures targets in the selection task can be selected if fixated at the center even if there is up to  $1^\circ$  of spatial error in the gaze signal [6].

$$X'_n = (\theta_n + N(0, \sigma), \psi_n + N(0, \sigma), 1)$$

**Temporal Downsampling:** Gaze positions were only sampled once every  $k = 30$  samples (500 ms) with a re-identification rate of 21.8% based on previous research [53].

$$\text{If } (n \bmod k = 0), \text{ then } (X'_n = X_n), \text{ else } (X'_n = X'_{n-1})$$

**Weighted Smoothing:** The weighted average of the past  $B = 50$  samples (833 ms) of raw gaze positions are used for the perturbed gaze position. Based on prior research, this results in the lowest re-identification rate (14.1%) while minimizing impact on utility [53].

$$X'_n = \frac{X_{n-B} + 2(X_{n+1-B}) + 3(X_{n+2-B}) + \dots + B(X_n)}{\sum_{i=1}^B (i)}$$

## 4 RESULTS

We present our study results across the dimensions of user experience (Sec. 4.1), privacy mechanism preferences (Sec. 4.2), attitudes towards eye-tracking data collection (Sec. 4.3), sentiments towards sharing eye-tracking data with different entities or purposes (Sec. 4.4), sharing privacy-enhanced data (Sec. 4.5), and preferences towards how VIs are deployed (Sec. 4.6).

For statistical analysis, we first tested each measure for normality using a Kolmogorov-Smirnov (K-S) test. The collected data were ordinal and captured through five-point Likert responses (presented on a scale from 0 to 4). All data were determined to not be normally distributed based on the K-S test ( $p < 0.05$ ), except for the UEQ data ( $p > 0.05$ ). Our testing relied on comparisons between two groups across one factor in all cases except the UEQ. The UEQ analysis was applied to four groups comparing post-trial data from raw gaze and the three privacy mechanism conditions. For all data besides the UEQ, we relied on non-parametric Wilcoxon signed-rank testing for paired comparisons and Wilcoxon rank-sum testing for unpaired samples on between-task comparisons. For UEQ analysis, we applied a one-way ANOVA across the four groups and further investigated dimensions with significant differences using post-hoc t-tests with Bonferroni correction. We also calculated effect size with a Friedman Test on eye-tracking data attitudes and data-sharing sentiments and report comparisons with a medium ( $0.3 < W < 0.5$ ) or large ( $W \geq 0.5$ ) effect [29].

### 4.1 User Experience with VIs

#### 4.1.1 Aggregate

Post-block analysis compared preferences between the VIs (Figure 4). Results were consistent with VR findings [43]: the tendril VI was seen as significantly more distracting ( $p < 0.05$ ), while also being significantly more useful ( $p < 0.001$ ), exciting ( $p < 0.05$ ), high quality ( $p < 0.05$ ), informative ( $p < 0.001$ ), and made users more aware of what objects they were looking at ( $p < 0.001$ ). Participants provided explanations for their ratings and noted the real-time feedback from the tendril helped them focus more precisely, specifically for Gaze Selection, which required accurate gaze pointing. Participants also made it clear they found the tendril interface to be distracting in open-ended prompts. One participant shared, “It was useful, but at times it made me




Fig. 4: Attitudes towards each VI. Significant differences are indicated by \* ( $p < 0.05$ ) and \*\*\* ( $p < 0.001$ ). Box and whisker plots show the quartiles, medians (lines), and averages (Xs).

too aware of where I was looking, which could get distracting.” Despite this, 63% of participants indicated a preference for the tendril interface in the post-experiment survey. The 22% of participants who preferred the eye icon justified their decision based on the lack of information displayed directly where they were looking.

#### 4.1.2 Between Task

Users who completed the Gaze Selection task preferred the tendril (83%) more than users who completed the Art Gallery (60%). Participants in the Gaze Selection task benefited from a visual indicator for gaze-based target selection. Art Gallery participants were prone to extraneous distraction from the tendril. Regardless, a majority of both groups found the tendril distracting but still preferred it for its ability to increase familiarity and understanding of eye-tracking data. For Gaze Selection, the icon interface was implemented within a context-relevant visual cue as opposed to the floating pair of eyeballs many found discomforting in VR [43]. We did not find different responses between the VIs by task, and we did not receive as strong of negative responses as in the VR setting, perhaps due to the visibility of the cue being reduced in AR when mixing light and colors from the real-world with the peripheral cue [25].

### 4.2 Privacy Mechanism Preferences

To better understand how participants perceived the different privacy mechanisms, we used the UEQ to measure 6 dimensions of user experience after each trial UEQ scores range from -3 (negative experience) to +3 (positive experience). To make the results more accessible, we present results with a translated scale range of 0 to 6. Scores above 3.8 are considered positive, indicating users generally liked or appreciated the aspect of the system. Scores between 2.2 and 3.8 are considered neutral, suggesting users had mixed feelings or were neither strongly positive nor negative. Scores below 2.2 are negative, reflecting dissatisfaction or challenges [34].




Fig. 5: Mean and standard error for UEQ dimensions across raw data and privacy mechanisms. Significant differences were found within Attractiveness, Perspicuity, Efficiency, and Dependability; and are indicated by \* ( $p < 0.05$ ), \*\* ( $p < 0.01$ ), and \*\*\* ( $p < .001$ ).

#### 4.2.1 Aggregate

Figure 5 presents the mean and standard deviations of UEQ scores across each dimension for all gaze data conditions and all participants, regardless of task. Gaussian scored the lowest on average. Smoothing performed best in all dimensions besides Novelty. Significant differences were identified within all dimensions besides Stimulation and Novelty, showing the largest differences between Raw Data and Gaussian Noise compared to Smoothing. Raw Data and Temporal Downsampling had no significant differences and were rated similarly. Qualitative feedback provided further insight.

**Smoothing:** Participants who preferred Smoothing emphasized its usability and lack of distraction. One participant specifically said, “Smoothing was less distracting and hurt my head a lot less than the others.” It was generally perceived as the most consistent and trustworthy privacy mechanism, with participants describing it as predictable and less intrusive. It scored the highest in Attractiveness (4.7), Perspicuity (5.1), and Efficiency (4.4), meaning users found it visually appealing, easy to understand, and effective for completing tasks. Novelty (3.4) was its weakest dimension, suggesting that while it was practical and enjoyable, users did not find it particularly innovative or exciting.

**Temporal Downsampling:** For Temporal Downsampling, participants appreciated that it provided minimal data exposure. As one participant commented, “Temporal Downsampling made me feel the safest because it shared the least amount of my eye data.” Two participants reported they felt it was distracting due to the latency it introduced. Participants who felt safest with this mechanism said it gathered the least amount of data “while still giving general area of data,” which was an important trade-off to them.

**Gaussian Noise:** In contrast, those who preferred Gaussian Noise viewed it as the most secure and the most challenging to use. One participant said, “The sporadic nature of the Gaussian Noise made it more difficult to pinpoint where I was looking, which made me feel more secure.” It was viewed as less reliable in usability but more secure for protecting user eye-tracking data. A participant explained, “I liked that the [Gaussian Noise] mechanism kept my eye gaze data scrambled so that you could never directly pinpoint the exact location I was looking at.” It received the lowest ratings for Attractiveness (2.9), Efficiency (3.0), Dependability (2.8), and Stimulation (3.6). Interestingly, it scored relatively well for Novelty as participants acknowledged its uniqueness in protecting privacy.




Fig. 6: Privacy Mechanism Preferences.

Participants were also asked to choose one of the three privacy mechanisms (or no preference) that made them feel the safest in terms of protecting their eye-tracking data at the end of the experiment. Figure 6 presents the distribution of preferences for all participants and by task. The majority of participants (50%) indicated they felt safest using Smoothing, followed by Gaussian Noise (24%) and Temporal Downsampling (19%). One participant in each group expressed no clear preference between the mechanisms.




Fig. 7: Mean and standard error for UEQ dimensions across Raw Data and privacy mechanisms in the Gaze Selection task. Significant difference was found between Attractiveness for Gaussian noise and Smoothing, indicated by \* ( $p < 0.05$ ).




Fig. 8: Mean and standard error for UEQ dimensions across Raw Data and privacy mechanisms in the Art Gallery task. Significant differences were found for Attractiveness, Perspicuity, and Dependability and are indicated by \* ( $p < 0.05$ ) and \*\* ( $p < 0.01$ ).

#### 4.2.2 Within Task

Figures 7 and 8 present the UEQ results by task. Similar trends were observed between the privacy mechanisms, with Smoothing scoring highest, followed by Temporal Downsampling, and then Gaussian noise, except for the Novelty dimension. The only significant difference within the Gaze Selection task data was the Attractiveness between Gaussian and Smoothing ( $p < 0.05$ ); the Art Gallery population had significant differences between mechanisms in Attractiveness, Perspicuity, and Dependability.

#### 4.2.3 Between Task

Figure 6 shows the Gaze Selection task population had a more balanced distribution of preferences for privacy mechanisms than the Art Gallery population which overwhelmingly preferred the Smoothing mechanism. We performed an independent samples t-test on the UEQ results between each task population. Significant differences were found in both the Temporal Downsampling and Smoothing mechanisms, specifically in the Attractiveness, Efficiency, and Dependability dimensions. The Attractiveness and Dependability scores were lower for Downsampling and Smoothing within the Gaze Selection task, while the Efficiency scores were lower in the Art Gallery task. This implies the smoother data streams from these mechanisms interrupted the Gaze-based Selection task less in comparison, while being less attractive in a more visually cluttered and dynamic environment.

### 4.3 Attitudes Towards Eye Tracking

Attitudes towards eye tracking and applications were measured on a five-point Likert scale (0: Strongly Disagree - 4: Strongly Agree).

#### 4.3.1 Aggregate

Participants’ familiarity with, understanding of, willingness to share (for benefits), and concerns (social, mental, physical, private) about eye-tracking data were collected before and after the experiment. Participants’ self-reported familiarity with and understanding of eye-tracking data both increased significantly ( $p < 0.001$ ) as the average ratings

(out of 4) increased from 2.4 to 3.0 for familiarity and 1.9 to 2.6 for understanding. Effect size was measured as medium ( $W = 0.44$ ) for familiarity and large ( $W = 0.5$ ) for understanding. Participants were less consistent about their willingness to share data for benefits, as the inner quartile ranges increased from 0.25 to 2, while the mean and median were similar.

Privacy concerns remained high both before and after the experiment, with 75% of participants rating Agree or Strongly Agree with the statement, “I am concerned about eye-tracking technology in terms of privacy.” When asked “I am concerned about eye-tracking technology in terms of social acceptability” before the experiment, social acceptability (Soc Concern) had the lowest average score (1.70: Neutral–Disagree). Privacy was the greatest concern with the highest average score (3.30: Agree–Strongly Agree), reflecting unease about data collection. Users reported learning more about eye-tracking data while the severity of their concerns (social, mental, physical, and private) remained similar.




Fig. 9: Attitudes towards eye-tracking data collection between pre and post-experiment surveys and by task. Box and whisker plots present the quartiles, medians (lines), and averages (Xs). Significant differences were found for Familiarity, Understanding, and Mental Concern and are indicated by a \* ( $p < 0.05$ ) and \*\*\* ( $p < 0.001$ ).

#### 4.3.2 Between Task

Participants in the Gaze Selection task had higher mental concerns ( $p < 0.05$ ), with a median of 3 (Agree) compared to a median of 2 (neutral). In addition, the change in mean rating of familiarity and understanding before and after the experiment for those who completed the Gaze Selection task was significant ( $p < 0.05$ ). All participants were more likely to report their perceived familiarity with and understanding of eye-tracking technology. The effect size was medium for both familiarity ( $W = 0.31$ ) and understanding ( $W = 0.44$ ) for those who experienced the Art Gallery task. A large effect size at  $W = 0.56$  was measured for the Gaze Selection task for both familiarity and understanding.

### 4.4 Data-Sharing Sentiments

Participants were asked about their willingness to share eye-tracking data before and after completing the experiment.

#### 4.4.1 Aggregate

Figures 10 and 11 present the data-sharing sentiments for specific purposes and entities, respectively. No significant differences were found between pre and post-experiment data-sharing sentiments.

Most participants disagreed with the idea of sharing their eye-tracking data to analyze shopping behavior or identify interests for targeted advertisements (medians of 1 or lower). However, participants were more receptive to sharing data for purposes that could benefit them directly, including enabling hands-free interaction (median: 3), improving user interfaces (median: 3), or monitoring stress levels (median: 4). In general, post-experiment scores were comparable or lower, except for tracking related to habits (activity tracking or lifelogging) which increased from neutral to slightly positive.

When asked about willingness to share data with specific entities or classifications of entities (general, government health agencies, etc.), the ratings between pre-experiment and post-experiment ratings remained generally the same. Additional variance was introduced in the post-experiment survey for sharing data with employers and to support features of XR applications. In general, participants were less willing to share data with entities that may be seen as holding the most power in their lives or private interests, such as the government, businesses, and employers. Overall, there were no significant differences in all data-sharing sentiments between the pre and post-experiment surveys.




Fig. 10: Data-sharing sentiments for purposes. Box and whisker plots present the quartiles, medians (lines), and averages (Xs).




Fig. 11: Data-sharing sentiments for entities. Box and whisker plots present the quartiles, medians (lines), and averages (Xs).

#### 4.4.2 Between Task

We compared responses for willingness to share data between the Gaze Selection and Art Gallery tasks and found no significant differences across purposes or entities.

### 4.5 Sharing Privacy-Enhanced Data

Participants were asked in the post-experiment survey to select any of the data conditions (Raw, Gaussian, Temporal Downsampling, Smoothing) they would be comfortable sharing. We found 47% of our participants still felt comfortable sharing Raw Data after experiencing VIs and privacy mechanisms. This is interesting given how strong responses were for concerns about eye-data privacy (Figure 9) from the post-block surveys. We hypothesize that while participants are concerned about privacy, the utility of gaze data outweighs it, as there is still a lack of awareness of how their data will be used to violate privacy for re-identification, ad targeting, and detecting sexual orientation or medical



conditions. The post-experiment data showed participants were more comfortable sharing data if privacy mechanisms were applied with 76% feeling comfortable with Gaussian Noise, 74% comfortable sharing data with Temporal Downsampling applied, and 70% with Smoothing.

## 4.6 Deployment of Interfaces

In the post-experiment survey, participants rated their agreement (from 0 to 4) recommending VIs be deployed in the following ways: used before running eye tracking, made available, enabled by default, and required by law. We found 64.7% of participants agreed VIs should be enabled by default. Furthermore, 94.1% of all participants agreed that VIs should be available and tried out before enabling eye-tracking applications. Generally, users agreed VIs be tried out before enabling eye tracking and available, with some agreement 70.5% they be legally required or enabled by default for applications using eye-tracking data.

Participants thoughts on VI deployment were collected at the end of each block. The post-block responses for whether the VI should be enabled by default or provided as an option are shown in Figure 4, with significant differences between VI type ( $p < 0.05$  and  $p < 0.001$ , respectively). This effect implies the interface type played a role in deployment preferences, with tendrils being preferred.

## 5 DISCUSSION

Our research goals were to explore how users perceive gaze-based VIs in AR applications and their preference for deployment. We evaluated VIs across two tasks to provide additional insights beyond free-viewing that was explored in existing work [43]. Additionally, we evaluated three privacy mechanisms to understand the trade-off between user experience and willingness to share data after privacy mechanisms are applied.

### 5.1 Value of Visceral Interfaces

Overall, we found 94.1% of participants ( $N = 32$ ) supported making VIs available, with 64.7% of participants ( $N = 22$ ) agreeing AR platforms should enable interfaces by default for all users. This diverges from existing findings in VR in which participants supported having access to VIs but indicated they should be controllable, with negative responses to them being enabled by default [43]. Participants were slightly less enthusiastic about requiring VIs to be provided by law, with 70% ( $N = 24$ ) agreeing. The general trend was that users value transparency and control when it comes to managing and understanding their eye-tracking data. The strong preference for pre-use deployment supports existing results that suggest VIs are best suited to on-boarding new users or educating policymakers or advocates who are new to eye-tracking technology [43].

Our findings also highlight the need to consider the expanded contextual and social dimensions of AR. The availability of privacy awareness for AR sensors has a stronger impact than niche VR use cases, as these devices have the potential to be worn and used continuously in everyday settings [40].

### 5.2 Privacy Mechanisms

Our study connects visceral notice with that of privacy mechanisms to inform users about the data-level privacy offered by mechanisms and experience the trade-off with data utility within the Gaze Selection task. We found that the privacy impact of mechanisms in quantitative studies did not necessarily correlate with what users found made them feel the most secure. For example, the same amount of participants rated feeling most secure using Temporal Downsampling and Gaussian Noise (21.2%) in the Post-Experiment survey, whereas prior research suggests Temporal Downsampling is much less effective at protecting privacy in the context of re-identification [20, 53].

We note the importance of understanding the relationships between interfaces and mechanisms within the context of AR applications. The ability to use VIs to experience privacy mechanisms provides the opportunity for co-design by developing new VIs and integrating subjective feedback from users while exploring new privacy mechanisms.

### 5.3 Implications on Data Sharing

An interesting result of our study was that nearly half (47%) of participants indicated they were comfortable sharing Raw Data in the post-experiment survey. This seems to contradict the data showing participants mostly agreed with having privacy concerns about eye-tracking data (Figure 9) and clear trends in their unwillingness to share data with certain entities (Figure 11). As indicated in Figure 10, participants see clear benefits to sharing eye-tracking data for certain applications, including features of XR applications and to support developers. When making a single decision to share data or not, participants felt the benefits outweighed the risks. While the presentation on privacy mechanisms explained possible privacy risks from sharing data, participants were not equipped with the context for how gaze data specifically could be used for these inferences and the current state-of-the-art performance of privacy attacks. Additional information about how the data is processed by applications is not made clear to the user. Participants were not equipped to be able to define a threshold for how much detail is too much. This result motivates us to explore broader data collection scenarios when deploying VIs and to develop novel VIs that leverage the visceral notice dimension of showing to relay information about how data is processed [12, 48].

### 5.4 Design Implications

**VI Availability:** Users agreed VIs should be implemented such that they should be available to AR users and tried before enabling eye tracking. Most users agreed VIs should be legally required to be available in AR apps and enabled in AR apps by default. This recommendation is consistent with the current data transparency policy for the Magic Leap 2 that requires a visual indicator to signify when applications collect any form of eye-tracking data [35]. Users wish to have access to VIs in their applications and control when the interfaces are enabled. In contrast, VR users indicated they wanted access to VIs but strongly indicated they did not want them enabled by default [43]. Open-sourced access will be made available for both AR VIs for future evaluations and AR deployments.

**User-Preferred VI:** There was a preference for the tendrils interface over the eye icon within and between tasks in an AR setting consistent with existing VR results [43]. Users expressed the tendrils was distracting but seemed more concerned with the lack of information relayed by the eye icons. Our AR adaptation for the Gaze Selection task made the icon interface less creepy by embodying it within a contextually relevant cat body, though this did not impact preference ratings of the interface. Users expressed the icons do not offer as much information as the tendrils interface, which was more important for Gaze Selection as it assisted in task performance. We note the tendrils could be made less distracting by highlighting fixations or modulating transparency based on dwell time on task-relevant objects or regions.

**Privacy Mechanisms:** Designing effective privacy mechanisms for AR requires solutions that balance privacy protection with utility. Smoothing was the most preferred mechanism, with high UEQ scores in all dimensions except Novelty. Gaussian noise provided a greater sense of security but significantly lower UEQ scores. Prior work has established that Weighted Smoothing outperforms Gaussian Noise in terms of lower re-identification rates [20, 53]. Additionally, Gaussian Noise is susceptible to mitigation by identifying the noise parameter and filtering the signal to reconstruct raw data while Smoothing is resistant to this kind of attack. This result suggests additional work is needed to align subjective perceptions of mechanisms with the quantitative impact on user privacy.

**Task Dependence:** User preferences for privacy mechanisms were influenced by the nature of the AR task and environment (Figures 6–8). These between-task differences show the type of activity influenced user preferences towards privacy mechanisms. For Gaze Selection, the Gaussian Noise with a standard deviation of  $1.5^\circ$  was not large enough to make the task more difficult, though if a larger amount of noise was used, we expect a strong shift in preferences based on the perceived utility of the data. Certain interfaces align better with certain tasks. Generally, tendrils dominated user preferences for both tasks. Participants in the Art Gallery task (60% Tendrils; 33% Icons; 7% no



preference) observed a higher appreciation for the icon interface than their Gaze Select peers (83% tendril; 17%).

## 5.5 Limitations

Our sample had diverse ethnicities, but may not fully represent the broader population, such as older age groups. People with different backgrounds may have unique challenges related to understanding AR technology, eye-tracking sensor data, and the role of privacy mechanisms in data sharing. Although the two AR tasks represent active and passive gaze applications, they do not cover the entire range of contexts where eye tracking is applied in AR, such as training, remote assistance, or collaboration. There may be discrepancies when comparing the data collected for the icon condition between tasks as the Gaze Selection task rendered the icon within the task-context using a cat compared to disembodied eyes in the Art Gallery. Mentioning privacy at the start of the study can lead to a priming effect on participants' survey responses. We administered pre-experiment surveys after explaining the role of privacy mechanisms to ensure any shift in attitudes was captured before experiencing VIs for direct comparison. The field of IoT and XR privacy and security has studies with [4, 45] and without [14, 15, 36] priming. These XR studies provide conflicting results on whether or not priming presents confounds including instances where participants realized the research groups' focus on privacy and security. Describing the role of privacy mechanisms was critical in our study to provide participants with a common understanding before the experiment, given our interest in evaluating the subjective perception of different privacy mechanism visualizations. Most of our results demonstrated small effect sizes; further research should be conducted to measure the long-term impact of experiencing visualizations of privacy-preserving mechanisms. Our participants were not experts in privacy or law. Therefore, the data should be viewed as a general takeaway from XR consumers.

## 5.6 Future Work

Future studies could incorporate more longitudinal data collections, such as observing how users interact with eye-tracking technology when VIs are present over long periods of time while using readily available applications. We are interested in studying user behavior in enabling or disabling VIs as they navigate new applications or environments over time and gain an increased understanding of their gaze data in different contexts. We see immense potential for these insights when considering nudging-based visceral interfaces for everyday AR [3]. There are clear steps to explore additional visualization approaches beyond the tendril and icon methods seen in existing work; including leveraging additional ideas from the initial discussion of VIs for XR gaze data [48] and in particular those tuned to provide notice about privacy mechanisms. For example, a slider tool could be integrated into the privacy mechanism selection process, allowing users to adjust the strength of privacy protection to record their preferred parameter values. We intend to explore individualized preferences to enable a deeper analysis of how human perceptions align with privacy-utility trade-offs.

Last, we are motivated by the finding that nearly half of the participants still felt comfortable sharing raw gaze data despite privacy concerns to develop novel VIs better suited to AR environments and explore the Showing dimension of visceral notice [12, 48]. Existing studies of gaze-based VIs neglected this dimension; we see an opportunity to develop VIs in the evolving space of everyday AR supported by contextual AI [16, 17]: the integration of Gemini AI with Google AR glasses [24] and the integration of low-power eye tracking into Sesame conversational AI glasses [33]. We envision a VI that surfaces notifications about current inferences from AI systems and what behavior led to them, as well as relaying information explaining how they are used by AI assistants [49, 54]. This type of interface would better inform users about how gaze data is linked to specific privacy risks, and it is critical as it fits the compelling use-case of ubiquitous head-worn AR devices.

## 6 CONCLUSION

AR technology has the potential to become a ubiquitous technology that enhances users' interactions at work, school, and home. However, the increased reliance on eye-tracking data to support critical applications has the potential to force the technology onto a large user base that does not understand the privacy implications. We explored two VIs for increasing privacy awareness towards eye-tracking data in AR settings and evaluated three privacy mechanisms to understand users' perceived privacy, comfort, and preferences. Users appreciated the utility of VIs and recommended that these tools be made available in AR apps. Our findings suggest clear privacy concerns for eye-tracking data and preferences for the Weighted Smoothing privacy mechanism. However, nearly half of our participants were still willing to share raw gaze data, seeming to weigh utility above privacy. This motivates a critical need to develop techniques capable of linking gaze data with the relevant privacy risks and establishing the standard use of privacy mechanisms. Our paper establishes connections between the implementation of VIs and privacy mechanisms and their potential to serve as tools that empower individual users to maintain autonomy over their privacy for emerging XR technologies.

## SUPPLEMENTAL MATERIALS

Our supplemental materials and codebase will be made publicly available at [https://bmdj-vt.github.io/project\\_pages/privacy\\_notice](https://bmdj-vt.github.io/project_pages/privacy_notice).

## ACKNOWLEDGMENTS

The authors acknowledge support from the National Science Foundation CSGGrad4US Fellowship (CNS-2240205) and the Commonwealth Cyber Initiative (CCI). The authors also thank Meagan Smith for her contributions in helping edit the manuscript.

## REFERENCES

- [1] M. Abraham, M. McGill, and M. Khamis. What you experience is what we collect: User experience based fine-grained permissions for everyday augmented reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–24, 2024. 1
- [2] M. Abraham, P. Saeghe, M. McGill, and M. Khamis. Implications of xr on privacy, security and behaviour: Insights from experts. In *Nordic Human-Computer Interaction Conference*, pp. 1–12, 2022. 1
- [3] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. F. Cranor, S. Komanduri, P. G. Leon, N. Sadeh, F. Schaub, M. Sleeper, et al. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3):1–41, 2017. 9
- [4] R. G. Anaraky, P. Bahirat, M. Nasiri, X. Page, B. P. Knijnenburg, and A. T. Duchowski. Effect of priming on smart home privacy preferences. 2020. 9
- [5] S. Aziz and O. Komogortsev. Assessing the privacy risk of cross-platform identity linkage using eye movement biometrics. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–9. IEEE, 2023. 1
- [6] S. Aziz, D. J. Lohr, L. Friedman, and O. Komogortsev. Evaluation of eye tracking signal quality for virtual reality applications: A case study in the meta quest pro. In *Proceedings of the 2024 Symposium on Eye Tracking Research and Applications*, pp. 1–8, 2024. 5
- [7] Y. Bar-Haim, T. Ziv, D. Lamy, and R. M. Hodes. Nature and nurture in own-race face processing. *Psychological science*, 17(2):159–163, 2006. 2
- [8] S. Bouchard, M. Berthiaume, G. Robillard, H. Forget, C. Daudelin-Peltier, P. Renaud, C. Blais, and D. Fiset. Arguing in favor of revising the simulator sickness questionnaire factor structure when assessing side effects induced by immersions in virtual reality. *Frontiers in Psychiatry*, 12:739742, 2021. 3
- [9] E. Bozkir, O. Günlü, W. Fuhl, R. F. Schaefer, and E. Kasneci. Differential privacy for eye tracking with temporal correlations. *Plos one*, 16(8):e0255979, 2021. 1, 2
- [10] E. Bozkir, S. Özdel, K. H. C. Lau, M. Wang, H. Gao, and E. Kasneci. Embedding large language models into extended reality: Opportunities and challenges for inclusion, engagement, and privacy. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces*, pp. 1–7, 2024. 1

- [11] E. Bozkir, S. Özdel, M. Wang, B. David-John, H. Gao, K. Butler, E. Jain, and E. Kasneci. Eye-tracked virtual reality: a comprehensive survey on methods and privacy challenges. *arXiv preprint arXiv:2305.14080*, 2023. 1
- [12] R. Calo. Against notice skepticism in privacy (and elsewhere). *Notre Dame L. Rev.*, 87:1027, 2011. 1, 2, 8, 9
- [13] Y.-H. Cha, J. F. Golding, B. Keshavarz, J. Furman, J.-S. Kim, J. A. Lopez-Escamez, M. Magnusson, B. J. Yates, B. D. Lawson, et al. Motion sickness diagnostic criteria: Consensus document of the classification committee of the bárány society. *Journal of Vestibular Research*, 31(5):327–344, 2021. 3
- [14] K. Cheng, A. Bhattacharya, M. Lin, J. Lee, A. Kumar, J. F. Tian, T. Kohno, and F. Roesner. When the user is inside the user interface: An empirical study of {UI} security properties in augmented reality. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 2707–2723, 2024. 9
- [15] K. Cheng, J. F. Tian, T. Kohno, and F. Roesner. Exploring user reactions and mental models towards perceptual manipulation attacks in mixed reality. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 911–928, 2023. 9
- [16] H. Cho, Y. Yan, K. Todi, M. Parent, M. Smith, T. R. Jonker, H. Benko, and D. Lindlbauer. Minexr: Mining personalized extended reality interfaces. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2024. 9
- [17] S. Davari, D. Stover, A. Giovannelli, C. Ilo, and D. A. Bowman. Towards intelligent augmented reality (iar): A taxonomy of context, an architecture for iar, and an empirical study. *arXiv preprint arXiv:2411.02684*, 2024. 9
- [18] B. David-John, K. Butler, and E. Jain. For your eyes only: Privacy-preserving eye-tracking datasets. In *2022 Symposium on Eye Tracking Research and Applications*, 2022. 1, 2
- [19] B. David-John, K. Butler, and E. Jain. Privacy-preserving datasets of eye-tracking samples with applications in xr. *IEEE Transactions on Visualization and Computer Graphics*, 29(5), 2023. 1, 2
- [20] B. David-John, D. Hosfelt, K. Butler, and E. Jain. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2555–2565, 2021. 1, 2, 8
- [21] J. A. De Guzman, K. Thilakarathna, and A. Seneviratne. Security and privacy approaches in mixed reality: A literature survey. *ACM Computing Surveys (CSUR)*, 52(6):1–37, 2019. 2
- [22] N. Deng, Z. He, J. Ye, B. Duinkharjav, P. Chakravarthula, X. Yang, and Q. Sun. Fov-nerf: Foveated neural radiance fields for virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3854–3864, 2022. 1
- [23] G. Diaz, J. Cooper, D. Kit, and M. Hayhoe. Real-time recording and classification of eye movements in an immersive virtual environment. *Journal of vision*, 13(12):5–5, 2013. 5
- [24] I. Fried. Google shows new ar glasses, vr headset at ted. <https://www.axios.com/2025/04/08/google-ar-glasses-vr-headset-ted>, 2024. [Accessed: 10-April-2025]. 9
- [25] J. L. Gabbard, M. Smith, C. Merenda, G. Burnett, and D. R. Large. A perceptual color-matching method for examining color blending in augmented reality head-up display graphics. *IEEE Transactions on Visualization and Computer Graphics*, 28(8):2834–2851, 2020. 5
- [26] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18995–19012, 2022. 2
- [27] S. G. Hart. NASA-Task Load Index (NASA-TLX); 20 Years Later. Technical report, NASA Ames Research Center, 2006. 3
- [28] B. Heller and A. Bar-Zeev. The problems with immersive advertising: in ar/vr, nobody knows you are an ad. *Journal of Online Trust and Safety*, 1(1), 2021. 2
- [29] A. Kassambara. Friedman Test Effect Size (Kendall's W Value). [https://rpkgstats.com/rstatix/reference/friedman\\_effsize.html](https://rpkgstats.com/rstatix/reference/friedman_effsize.html), 2022. 5
- [30] C. Katsini, Y. Abdrabou, G. E. Raptis, M. Khamis, and F. Alt. The role of eye gaze in security and privacy applications: Survey and future hci research directions. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pp. 1–21, 2020. 1
- [31] Y. Kim. Virtual reality data and its privacy regulatory challenges: A call to move beyond text-based informed consent. *Cal. L. Rev.*, 110:225, 2022. 2
- [32] J. L. Kröger, O. H.-M. Lutz, and F. Müller. What does your gaze reveal about you? on the privacy implications of eye tracking. In *IFIP International Summer School on Privacy and Identity Management*, pp. 226–241. Springer, 2020. 1, 2
- [33] Z. Labs. Zinn labs and sesame. <https://www.zinnlabs.com/>, 2025. [Accessed: 20-April-2025]. 9
- [34] B. Laugwitz, T. Held, and M. Schrepp. Construction and evaluation of a user experience questionnaire. In *HCI and Usability for Education and Work: 4th Symposium of the Workgroup Human-Computer Interaction and Usability Engineering of the Austrian Computer Society, USAB 2008, Graz, Austria, November 20-21, 2008. Proceedings 4*, pp. 63–76. Springer, 2008. 4, 5
- [35] M. Leap. Magic leap 2 eye tracking data transparency policy. <https://www.magicleap.com/legal/eye-tracking>, 2023. [Accessed: 25-April-2024]. 2, 8
- [36] J. Li, A. R. Chowdhury, K. Fawaz, and Y. Kim. {Kaleido}:{Real-Time} privacy control for {Eye-Tracking} systems. In *30th USENIX security symposium (USENIX security 21)*, pp. 1793–1810, 2021. 2, 9
- [37] J. Lin, J. Cronjé, C. Wienrich, P. Pauli, and M. E. Latoschik. Visual indicators representing avatars' authenticity in social virtual reality and their impacts on perceived trustworthiness. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
- [38] D. Lohr, S. Aziz, L. Friedman, and O. V. Komogortsev. Gazebasevr, a large-scale, longitudinal, binocular eye-tracking dataset collected in virtual reality. *Scientific Data*, 10(1):177, 2023. 1
- [39] D. Lohr, H. Griffith, and O. V. Komogortsev. Eye know you: Metric learning for end-to-end biometric authentication using eye movements from a longitudinal dataset. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(2):276–288, 2022. 1
- [40] F. Lu and D. A. Bowman. Evaluating the potential of glanceable ar interfaces for authentic everyday uses. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 768–777. IEEE, 2021. 8
- [41] D. Munoz, J. Broughton, J. Goldring, and I. Armstrong. Age-related performance of human subjects on saccadic eye movement tasks. *Experimental brain research*, 121:391–400, 1998. 2
- [42] K. Pfeuffer, B. Mayer, D. Mardanbegi, and H. Gellersen. Gaze+ pinch interaction in virtual reality. In *Proceedings of the 5th symposium on spatial user interaction*, pp. 99–108, 2017. 1
- [43] G. N. Ramirez-Saffy, P. K. Chelladurai, A. Vargas, I. Bukhari, E. Selinger, S. Foster, B. Heller, and B. David-John. Visceral interfaces for privacy awareness of eye tracking in vr. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 396–405. IEEE, 2024. 1, 2, 3, 4, 5, 8
- [44] A. ROSSI. *Legal Design for the General Data Protection Regulation. A Methodology for the Visualization and Communication of Legal Concepts*. PhD thesis, UNIBO - Università di Bologna, Bologna, Italy, 29 March 2019. 2
- [45] M. Sajid, S. I. M. Shah Bukhari, B. Ji, and B. David-John. "Just stop doing everything for now!": Understanding security attacks in remote collaborative mixed reality. In *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 623–633. IEEE Computer Society, Los Alamitos, CA, USA, Mar. 2025. doi: 10.1109/VR59515.2025.00085 9
- [46] N. Sammaknejad, H. Pouretamad, C. Eslahchi, A. Salahirad, and A. Alinejad. Gender classification based on eye movements: A processing effect during passive face viewing. *Advances in cognitive psychology*, 13(3):232, 2017. 2
- [47] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A design space for effective privacy notices. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*, pp. 1–17, 2015. 2
- [48] E. Selinger, E. Altman, and S. Foster. Eye-tracking in virtual reality: A visceral notice approach for protecting privacy. *Privacy Studies Journal*, 2:1–34, Mar. 2023. doi: 10.7146/psj.v2i.134656 1, 2, 4, 8, 9
- [49] N. Sendhilnathan, A. S. Fernandes, M. J. Proulx, and T. R. Jonker. Implicit gaze research for xr systems. *arXiv preprint arXiv:2405.13878*, 2024. 9
- [50] R. Singh, M. Huzaifa, J. Liu, A. Patney, H. Sharif, Y. Zhao, and S. Adve. Power, performance, and image quality tradeoffs in foveated rendering. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 205–214. IEEE, 2023. 1
- [51] J. Steil, I. Hagedstedt, M. X. Huang, and A. Bulling. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pp. 1–9, 2019. 1, 2
- [52] A. E. Waldman. Privacy, notice, and design. *Stan. Tech. L. Rev.*, 21:74, 2018. 2
- [53] E. Wilson, A. Ibragimov, M. J. Proulx, S. D. Tetali, K. Butler, and E. Jain. Privacy-preserving gaze data streaming in immersive interactive virtual re-

ality: Robustness and user experience. *IEEE Transactions on Visualization and Computer Graphics*, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)

- [54] E. Wilson, N. Sendhilnathan, C. S. Burlingham, Y. Mansour, R. Cavin, S. D. Tetali, A. S. Fernandes, and M. J. Proulx. Eye gaze as a signal for conveying user attention in contextual ai systems. *arXiv preprint arXiv:2501.13878*, 2025. [9](#)
- [55] M. Windl, A. Scheidle, C. George, and S. Mayer. Investigating security indicators for hyperlinking within the metaverse. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*, 2023. [2](#)
- [56] M. Yu, D. Harris, I. Jones, T. Zhang, Y. Liu, N. Sendhilnathan, N. Kokhlikyan, F. Wang, C. Tran, J. L. Livingston, et al. Explainable interfaces for rapid gaze-based interactions in mixed reality. *arXiv preprint arXiv:2404.13777*, 2024. [1](#)