JDS

Visceral Notices: Rethinking Consent for Passive Sensing in Augmented Reality

NISSI OTOO, Virginia Tech, USA
G. NIKKI RAMIREZ, Virginia Tech, USA
EVAN SELINGER, Rochester Institute of Technology, USA
SHAUN FOSTER, Rochester Institute of Technology, USA
BRENDAN DAVID-JOHN, Virginia Tech, USA

Introduction

Augmented reality (AR) headsets are transitioning from specialized tools to everyday devices, especially in immersive (VR/AR) applications. Eye tracking, essential for foveated rendering, intuitive interaction, and contextual AI assistance, has become more widely used in commercial AR platforms [9]. Unlike explicit user actions (e.g., clicks, voice commands), gaze is reactive and cannot be consciously controlled. This creates a unique consent challenge. Eye-tracking data reveals intimate, sensitive information through inference: medical conditions, sexual orientation, cognitive states, personality traits, substance use, and emotional vulnerabilities [7]. As Kroger et al. observe, individuals cannot opt out of revealing information through their gaze patterns, even when they are aware of potential privacy risks [7].

Our team has researched visceral interfaces, immersive VR/AR visualizations of gaze data building on the concept of visceral privacy notice [8, 10, 12]. However, recent studies suggest there is still a gap in users' protective data sharing decisions. If experiential, spatial, real-time visualization, the most immersive form of notice possible, still fails to translate privacy awareness into protective behavior, what does this imply for meaningful consent? In this position paper, we argue for consequence-showing interfaces that visualize what gaze data reveals, stakes-based default protections in high-risk contexts, and consent mechanisms triggered at inference generation rather than data collection.

Visceral Notice for VR/AR Eye Tracking

Calo introduced visceral notice as an experientially rich approach to privacy awareness that leverages psychological and sensory cues to communicate data collection practices [4]. Unlike traditional text-based notices, visceral interfaces make privacy implications tangible through familiar stimuli that elicit immediate understanding. This concept was applied to VR eye tracking, proposing interfaces that leverage "familiarity as warning" (e.g., animated eyes watching the user) and "showing" (displaying gaze trails in real time) [12]. Ramirez-Saffy et al. conducted the first empirical evaluation of these interfaces in VR, finding that tendril visualizations showing real-time gaze trails significantly increased privacy awareness and reduced data-sharing willingness in free-viewing tasks [10]. Our prior work extends these findings to AR contexts with both passive (art gallery) and active (gaze-based selection) tasks [8].

Context-Dependent Consent

How does the threshold for being informed change depending on context? Our position is that for AR sensor data it should not, at least not in the way current consent frameworks assume. Users cannot reliably assess stakes for emerging behavioral sensing. Otoo et al.'s study provides evidence of this failure [8]. Despite experiencing visceral interfaces that made gaze collection experientially real, 47% of participants still felt comfortable sharing raw eye-tracking data. When asked further about inferences

Authors' Contact Information: Nissi Otoo, nissiotoo@vt.edu, Virginia Tech, USA; G. Nikki Ramirez, gnramirezsaffy@vt.edu, Virginia Tech, USA; Evan Selinger, eselinger@gmail.com, Rochester Institute of Technology, USA; Shaun Foster, scffaa@rit.edu, Rochester Institute of Technology, USA; Brendan David-John, bmdj@vt.edu, Virginia Tech, USA.

made on gaze data, 75% of participants indicated clear privacy concerns. The problem is that their concern does not translate to protective behavior when the stakes are not immediately obvious.

This awareness-protection gap becomes dangerous in high-stakes contexts that users do not recognize as such. For example, pupil dilation patterns and eye movements reveal stress levels [7]. Users knows their gaze is being tracked since they saw the visceral interface during onboarding, but they don't know a potential employer "productivity dashboard" is receiving stress metrics derived from that data. From the user's perspective, it's a productivity tool, but from a privacy perspective, it's workplace surveillance. The same dynamic plays out in public spaces. Roesner et al. identified facial recognition in AR as a primary bystander privacy risk years ago [11], yet the threat risk has expanded with modern AR smart glasses. Modern AR glasses perform continuous scene understanding, and gaze data reveals not just what you looked at, but patterns that enable re-identification across contexts [3, 5].

Healthcare contexts present similar problems. Eye-tracking metrics can indicate neurological conditions, substance use, and cognitive decline [7]. A user trying an AR meditation app sees a visceral interface showing their gaze trails and thinks this is just wellness tracking. They do not consider that gaze patterns indicating early-stage dementia might flow to insurance risk models [6]. The context feels low-stakes because the interface shows data collection, not data use. This aligns with AI regulation frameworks that distinguish low-, medium-, and high-risk AI applications. The EU AI Act establishes that higher-risk scenarios demand proportionally stronger safeguards, and that policymakers, not individual users, should classify risk levels [2]. Therefore, this motivates a design principle we see as essential for AR: high-stakes contexts need default protections, not just better notice. Additionally, there is a need to explore new visceral interfaces that are linked directly to inferences from shared data.

Collection Notice to Inference Consequence

How might AI systems introduce new forms of nudging or manipulation? We also propose a new category of visceral interface: real-time inference visualization. Instead of showing data collection, these interfaces show data use. For example, a user receives an AR notification that surfaces when an AI system generates a sensitive inference: *"Stress level: High (detected from pupil dilation + fixation instability). Data sent to: employer productivity dashboard, updated 3 minutes ago."* or *"Sexual orientation estimate updated (confidence: 73%, based on gaze dwell patterns on faces). Currently shared with: ad network serving dating app ads."* This applies friction at the moment of inference generation, not just data collection. The distinction matters because users do not experience privacy loss when their gaze is tracked. They experience it when that tracking reveals information they consider private. Traditional visceral interfaces show the cause. Inference visualization would show the effect. We note that AI compounds this challenge by making weak or passive signals productive. Even low-quality gaze data becomes meaningful when AI layers extract features that enable inference, widening the gap between what users think they're consenting to and what systems can infer.

If we design interfaces that nudge users toward protective behavior by surfacing alarming inferences, are we manipulating them? Setting defaults is both inevitable and ethical. Every interface design choice nudges users toward some behavior. The ethical question is not whether to nudge, but whether the nudge serves user interests. In high-risk scenarios where users demonstrably cannot protect themselves despite understanding surveillance, nudging toward protection is justified. However, this introduces an ethical tension. Acquisti et al. argue that effective privacy nudges should enhance user autonomy rather than replace it [1]. Our position is nudging toward protection is only justified

when the alternative is exploitation through ignorance. For example, users who do not understand that their gaze patterns reveal sexual orientation or cognitive decline are not exercising informed consent.

Conclusion

We argue for a fundamental reframing that consent should occur at inference generation, not data collection. High-stakes contexts such as employment monitoring, public surveillance, and medical inference require mandatory privacy mechanisms. Additionally, visceral interfaces must evolve beyond showing what is collected to showing what is revealed. We bring questions into the workshop such as, "*What if AR systems surfaced real-time notifications when sensitive inferences are generated*?", "*What if users could contest AI predictions about their cognitive state, sexual orientation, or stress levels?*", "*What if friction occurred not when you agree to tracking, but when tracking enables decisions about you?*" Collection-based consent fails because users cannot protect themselves from inferences they cannot prevent through eye movements they cannot suppress.

Author Bios

Nissi Otoo is an undergraduate student at Virginia Tech whose work examines user understanding in mixed reality, with a focus in eye tracking, and technology equity. G. Nikki Ramirez is a PhD student at Virginia Tech. Her research focuses on deceptive risks & perceptual manipulation in virtual reality. Evan Selinger is a Professor of Philosophy at Rochester Institute of Technology, and his research focuses on the ethical & legal dimensions of technology. Shaun Foster is a Professor of 3D Digital Design at Rochester Institute of Technology, an RIT Generative AI Faculty Fellow, and an Authorized Unreal Instructor whose research connects national broadcast, advertising, & immersive education, with current focus on AI-driven 3D workflows, XR/VR eye tracking. Brendan David-John is an Assistant Professor in Computer Science at Virginia Tech with a focus on XR privacy & security, adaptive interfaces, and eye tracking.

References

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–41.

[2] EU Artificial Intelligence Act. 2024. The eu artificial intelligence act. *European Union* (2024).

[3] Samantha Aziz and Oleg Komogortsev. 2023. Assessing the privacy risk of cross-platform identity linkage using eye movement biometrics. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 1–9.

[4] Ryan Calo. 2011. Against notice skepticism in privacy (and elsewhere). *Notre Dame L. Rev.* 87 (2011), 1027.

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

[6] Brittan Heller and Avi Bar-Zeev. 2021. The problems with immersive advertising: in AR/VR, nobody knows you are an ad. *Journal of Online Trust and Safety* 1, 1 (2021).

[7] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What does your gaze reveal about you? On the privacy implications of eye tracking. In *IFIP International Summer School on Privacy and Identity Management*. Springer, 226–241.

[8] Nissi Otoo, Kailon Blue, G Nikki Ramirez, Evan Selinger, Shaun Foster, and Brendan David-John. 2025. Visceral notices and privacy mechanisms for eye tracking in augmented reality. *IEEE Transactions on Visualization and Computer Graphics* (2025).

[9] Ken Pfeuffer, Benedikt Mayer, Diako Mardanbegi, and Hans Gellersen. 2017. Gaze+ pinch interaction in virtual reality. In *Proceedings of the 5th symposium on spatial user interaction*. 99–108.

[10] G Nikki Ramirez-Saffy, Pratheep Kumar Chelladurai, Alances Vargas, Ibrahim Bukhari, Evan Selinger, Shaun Foster, Brittan Heller, and Brendan David-John. 2024. Visceral Interfaces for Privacy Awareness of Eye Tracking in VR. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 396–405.

[11] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (2014), 88–96.

[12] Evan Selinger, Ely Altman, and Shaun Foster. 2023. Eye-Tracking in Virtual Reality: A Visceral Notice Approach for Protecting Privacy. *Privacy Studies Journal* 2 (Mar. 2023), 1–34. doi:10.7146/psj.v2i.134656