Let's SOUP up XR: Collected thoughts from an IEEE VR workshop on privacy in mixed reality

Brendan David-John
University of Florida

Diane Hosfelt
Independent Researcher

Kevin Butler
University of Florida

Eakta Jain
University of Florida

Abstract

This paper presents insights from the PrXR workshop conducted at IEEE VR 2021. We identified several topic areas related to privacy and security risks for virtual, augmented, and mixed-reality (XR) applications. Risks are presented from the perspective of the XR community. We attempt to thematically group the workshop findings and highlight the challenges brought up by the participants. The identified research topics serve as a roadmap to push forward privacy and security research in the context of XR.

1 Introduction

As mixed-reality devices become more pervasive in everyday life,[1] it is imperative that researchers consider the myriad of privacy and security challenges presented by these devices. Huge amounts of data are being collected and stored with the consequences of a data leak being opaque to most users. Informed users are indeed concerned about these risks, and rightfully so based on what can be revealed [2]. Beyond the risk of inferring information that was thought to be private, keeping this data secure at such a massive scale is a corresponding security challenge.

The PrXR workshop was conducted to take advantage of momentum from a keynote at IEEE VR 2020 by Dr. Susan Persky that raised many concerns about data privacy in the context of medical applications and inferences.[2] Privacy discussions were a recurring theme in 2021 as part of Dr. Steven Feiner's keynote and Dr. David Luebke's award acceptance speech.[3,4] The organizers proposed the PrXR workshop in an effort to gather the perspectives of the XR community.

The IEEE VR conference used the Virbela platform, which inspired attendees' discussion. In Virbela, each attendee was represented by a virtual avatar which they could customize by selecting various geometries and personalize the color of the skin, hair, and clothes. Within the virtual environment, privacy circles were marked on the ground that provided a 'cone of silence' for users, ensuring their conversation was not heard outside of the region. The characteristics of the virtual environment grounded the workshop discussions with a concrete example of a virtual world.

This paper is organized as follows: Section 2 summarizes insights from our paper presentations and panel of invited experts, Section 3 collects breakout room discussions seeded by risk categories defined by the organizers, and Section 4 connects the landscape of privacy and security research as seen by XR researchers with the recent work in this direction within the USENIX community.

2 Position paper and panelist perspectives

Among the position paper presentations and panelist discussions were several highlights that are key to both the XR and privacy research communities. First, the panelists discussed the responsibilities of non-privacy researchers in the XR community. Next, they described the differences introduced by a shift from 2D computing to immersive 3D environments and then identified existing threats to consider and current approaches to addressing them.

---

[1]Consider that many phones are capable of augmented reality and the Oculus Quest 2 sold two to three million units in Q4 of 2020 [15] which is more than all other Oculus devices combined.

---

[2]https://ieeevr.org/2020/program/speakers.html#persky
[3]https://ieeevr.org/2021/program/keynote-speakers/#keynote-feiner
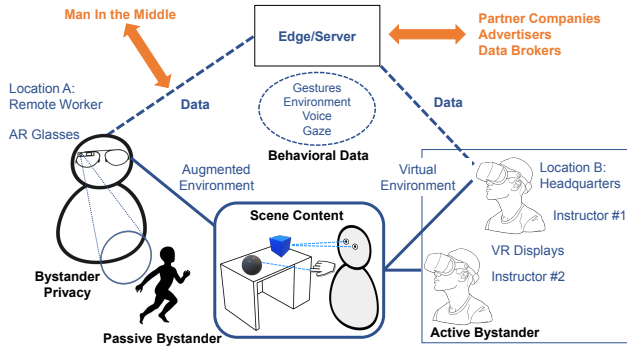[4]https://ieeevr.org/2021/awards/vgtc-award-winners/

Figure 1: Overview of data collected and streamed in an XR scenario. Shared environments rely on sharing data with edge devices or servers for processing, which introduces the risk of man-in-the-middle attacks and enables sharing of aggregated data with third parties that may act as bad actors. The assets within this figure are available by Creative Commons licenses [13, 34, 36]

2.1 Researcher Responsibilities

XR researchers and developers naturally look toward the novel use cases and associated quality of experience advantages created by advances in hardware, ease of content creation, and market penetration. However, the field has now reached a level of maturity where it is necessary to consider the balance between technological benefits and the resulting privacy and security risks. Happa and colleagues proposed that certification and standards are required at this stage to assess the risks and benefits of proposed XR technology and enable safe deployment through privacy benchmarks [14].

Panelists proposed that teaching ethical responsibilities is a necessary component in curricula for XR computing and HCI research, much like the Hippocratic oath exists to prevent harm resulting from doctors in the medical field. The challenges of introducing ethics in a computing curriculum include perceptions by students and faculty that others outside of computing should address them. This passing of the proverbial baton to either law makers or practitioners is based on the assumption that companies and platforms will collect data, simply because they can. A future in which technology is quantified in terms of risk during design and development is the key to exploring new technologies while also protecting users of the technology.

2.2 Challenges created by immersion

XR devices create virtual content that the user perceives as occupying the space around them. Buck and Bodenheimer presented privacy considerations related to the user's personal space [5]. Panelists and attendees additionally brought up adversarial scenarios such as changing what visitors see when they visit your home, or creating a negative first impression

during a job interview by changing the user's appearance. These discussions likely drew upon attendees' experiences in customizing their avatar and conjectures around what a virtual world platform, such as the one the conference was using, could change at will without consulting the attendees.

2.3 Risks from Large-scale Data Collection

One of the most pressing risks for the future of XR is related to data privacy at a large scale and the types of private information that can be inferred from machine learning models. On the one hand, machine learning is key to enabling many of the useful applications that are promised by mixed-reality devices, such as spatial mapping or naturalistic interfaces; but on the other hand, we are aggregating data that has the unknown potential to infer data that was thought to be private [28]. It is assumed that large amounts of individual user data will be captured by XR devices then collected and processed by edge devices or cloud servers, as illustrated in Figure 1. Currently there is no framework to ensure data privacy when actively aggregating data from multiple sensors at a massive scale, ranging from gaze and audio data to the way users interact with digital or physical objects. With the potential to reveal sensitive data about user's behavior and environment we have to consider best practices in protecting data privacy prior to risks being identified or users being harmed. Existing technical solutions include differential privacy and privacy-preserving federated learning, which have been applied to protect individuals within aggregated datasets and trained models [4, 11, 23, 26, 33, 35, 37].

3 Risk Categories

The breakout discussions were separated into three thematic areas for privacy and security considerations created by large-scale XR deployment (Figure 1). Each group was presented with the Gatekeeper model of David-John et al. [8] as a starting point for discussion (Figure 2). Some breakout groups focused on uses of the Gatekeeper for protecting data sources, while others used the data streaming model as a thought exercise to approach how data flows through an XR system.

3.1 Behavioral Data

Because XR devices require a variety of data about the user for their functionality, an app or third party affiliate can make a number of inferences about the user from the data streams. Undesired inferences could include correlating multiple data sources to de-anonymize users. An example famous to the SOUPS community is that of Narayanan and Shmatikov, who demonstrated that anonymous users of Netflix could be correlated with just a small amount of additional information from the Internet Movie Database in order to easily identify individual subscribers [27].

Figure 2: Gatekeeper model for enabling gaze applications without sharing raw gaze sample data [8]. This example shows how data from a single sensor can be processed in a privacy-preserving manner and provides a foundation for developing threat models and scenarios for XR data.

While mobile devices collect some of the same types of data, for example, accelerometer data, XR is going to involve high-fidelity data collection of eye-tracking data, motion-tracking data, heart rate data, and muscle activation data. As shown in Figure 1, data from multiple sources may be aggregated and processed on an edge device or cloud server.

Moreover, while controlling access to individual sensors could be a first step (e.g., to remove the risk of biometric identification), the fusion of data sensors is not currently regulated by law, nor is it clear how this would be accomplished. Heller introduces the notion of "biometric psychography", the practice of using biometric data to identify a person's interest through behavioral and anatomical information and discusses the three states that currently have laws relating to biometric privacy [16]. For example, while the Gatekeeper model shown in Figure 2 may protect against biometric identification from gaze data, data aggregators may use the Area-of-Interest statistics to generate personalized ads or infer sensitive characteristics like sexual preference [29]. A larger framework for assessing such risks, particularly in the context of user experience, are still an open area of investigation.

Beyond this low-level captured data, another important factor is how data is used to inform the design of the virtual space. Privacy can be relational and contextual; there are times when interaction is welcome, and times when one wants to be alone. Notably, these considerations can have implications on the representation of avatars in virtual environments.

A potential way forward for assuring behavioral privacy is making use of privacy indicators. For example, PhishAR is a startup that uses AR as a means to detect whether social engineering attacks such as phishing are being employed against a user [1]. Such indicators could also be transferred into the XR domain for other applications in order for users to assess the amount of private information being transmitted by the underlying platform, with potential settings in order to allow the users better control of how their data is used. The notion of "tangible privacy" was shared at the workshop [3]. Another instance discussed at the workshop was how the Virbela platform used for the main conference provided some privacy-preserving features in the form of visually marked regions that supported private conversations. However, atten-

dees noted that at the same the application was still collecting large amounts of data from users with potential risks, including the location of each user throughout the conference and who they interacted and had private conversations with. In this scenario we were trusting Virbela to maintain privacy. Future research should explore how to implement privacy methods for different forms of behavioral data and design of privacy-preserving XR platforms and interfaces.

3.2 Scene Content

A scene is composed of various types of content, such as the environment surrounding the user as well as feedback from the user themselves. Additionally, we consider the digital environment to be part of the scene content. As shown in Figure 1, remote workers can interact with a shared scene through both AR and VR displays. All of this information is aggregated and analyzed by edge devices or cloud servers. With such a broad context, it is clear that there are equally broad privacy considerations.

This breakout discussion focused on the physical scene in an AR context and culminated in an attack scenario wherein an application can sense the physical world and overlay virtual objects on it. For example, vehicle manufacturers are investigating in-car usage of AR [12]. Additionally, they have been looking at ways to monetize the in-car experience, so it is reasonable to expect that advertisements may be arriving to the driving experience [21]. If an ad is incorrectly or maliciously placed, it could occlude traffic information such as a stop sign, causing danger to drivers. As discussed previously by Roesner et al., it is more difficult for users to distinguish between real and virtual information in this paradigm [30]. Mitigations for risks related to scene content include having trusted offline data, minimizing the data given to applications, and having a guardian boundary that blends virtual content when you get close to physical objects [7].

Since scene content is such a large category, it needs further exploration beyond the consideration of physical space and occlusion. Next steps may explore how to protect interactions with the digital environment or investigating other aspects of the physical space.

3.3 Bystander Privacy

Threats to bystander privacy in XR are heightened due to the sheer number of sensors on mixed-reality devices that are expected to be worn all day long in public spaces. Mixed-reality devices contain motion and depth sensors that are used to model and reconstruct the environment, and might also attempt to classify the gender of bystanders or track movements to identify those with disabilities.

Mixed-reality displays are unique in that they could potentially modulate the appearance of bystanders, such as lightening skin tone or swapping professional business clothes for an

unprofessional or inappropriate appearance without consent.

If a passive bystander is not be able to provide consent to the XR device, the Gatekeeper (Figure 2) could remove the bystanders from the data passed along to the apps, or place virtual 'stickers' and emojis over their faces as has been previously proposed in the wearable camera context by Korayem et al. [20]. A Gatekeeper model could also support active bystanders, in which active implies that individual is able to interface with the privacy policy of nearby mixed-reality sensors. Interfacing with an active bystander can be achieved through a privacy token or their own mixed-reality device. An active bystander could then control what metadata is tracked and utilized by other devices, such as ethnicity or gender demographics, and what data is allowed to be manipulated about their appearance, if any. Figure 1 illustrates a passive bystander as a child who cannot provide consent while being recorded, and an active bystander as Instructor #2 who is being tracked by external cameras on the VR display of Instructor #1. Ruth et al. provide a scenario in which a virtual 'kick me' sign could be attached to a bystander that is being tracked by a mixed-reality user, however addressing this risk was outside of the scope of their proposed security solution for multiple users [32]. Privacy solutions that support active bystanders in the XR ecosystem is an open area of research that could be explored with existing devices.

4 Conclusion

This workshop aimed to foster dialogue around the privacy problems in XR by bringing expert voices from the security and privacy community to researchers and practitioners gathered for the largest academic conference in virtual/augmented/mixed reality. There is an emerging concern in this community about the risks created when physical reality and virtual reality intermix seamlessly. While previously it had been said that XR devices are simply another kind of mobile device, there is an increasing realization that there are unique threats in this space, ranging from novel attacks by malicious apps to novel inferences that can be made about the user. In the following paragraphs, we ruminate on what is different about XR from the perspective of XR researchers. We give some examples of works from the security and privacy (S&P) community that relate to this area. For the S&P perspective, we direct readers to an exhaustive survey by De Guzman and colleagues [9].

Spatial sensors are in the critical path of device operation Smartphones, tablets and smart wearables contain cameras, accelerometers, and location sensors. While a mobile phone can still be used for calling and texting if the user turns off location sensing, a mixed-reality 'call' relies on spatial sensing to appropriately render their conversation partner. In other words, certain sensors that are optional in the smartphone domain are operation-critical in the XR domain. Approaches such as world-driven access control [31] and feature sanitization [17] would turn off these sensors in agreed-upon private

environments, and are relevant for areas such as public bathrooms. Approaches such as plane releasing [10] are emerging solutions designed with XR operations in mind.

Novel sensors that track the user Users' body movements, gestures, head pose, and eye movements are tracked to a degree that was not required by mobile smartphones. Our work has exposed threats created by these novel sensors, for example, the user's iris biometric being exposed as a by-product of current eye tracking data flow [19]. Eye tracking data, in particular, can reveal intimate and personal detail about a user and is now being discussed in the S&P community [22]. Cross-pollination through workshops such as these would bring formal guarantees to the XR community and novel use cases to the S&P community. For example, the Kaleido system evaluates utility for gaze-as-pointer in a shooter game. It would be interesting to consider how the same system balances biometric identification with utility when data is used to animate the eyes of a virtual conversational avatar, as in [18].

Perceptual judgements can be altered Research has found that changing the virtual height of the user or the size of their arm can alter their perception of distance. Buck et al. showed that in VR the perceived environment and social context play a significant role in a user's personal space [6]. Manipulating personal space during interaction can influence user perception of safety and be used to expose biases or preferences. Furthermore, the ability to alter perception during locomotion is dangerous as it allows an adversary to influence unconscious behaviors of an individual and cause physical harm [7].

Avatars that represent us Conversations in mixed reality are going to rely on avatars that represent the user to their conversation partner. Induced privacy threats include longitudinal data on attributes such as how a user dresses their avatar. Novel attacks include manipulations to a user's avatar without their consent. Maloney et al. discusses the current landscape of avatar perception for social VR and discuss privacy implications and ethical guidelines [24, 25]. Key literature on how avatars are perceived could inform the S&P community on the impact of attack vectors and novel threats.

References

[1] PhishAR. https://www.phishar.com. Accessed: 2021-05-26.

[2] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M Redmiles. Ethics emerging: the story of privacy and security perceptions in virtual reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pages 427–442, 2018.

[3] Imtiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. Tangible privacy: Towards user-centric sensor designs for bystander privacy. *Proceedings of the ACM Journal: Human-Computer Interaction: Computer Supported Cooperative Work and Social Computing (CSCW '20)*, 4(CSCW2):116:1–116:28, October 2020.

[4] Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F Schaefer, and Enkelejda Kasneci. Differential privacy for eye tracking with temporal correlations. *arXiv preprint arXiv:2002.08972*, 2020.

[5] Lauren E Buck and Bobby Bodenheimer. Privacy and personal space: Addressing interactions and interaction data as a privacy concern. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 399–400. IEEE, 2021.

[6] Lauren E Buck, Sohee Park, and Bobby Bodenheimer. Determining peripersonal space boundaries and their plasticity in relation to object and agent characteristics in an immersive virtual environment. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 332–342. IEEE, 2020.

[7] Peter Casey, Ibrahim Baggili, and Ananya Yarramreddy. Immersive virtual reality attacks and the human joystick. *IEEE Transactions on Dependable and Secure Computing*, 2019.

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

[9] 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.

[10] Jaybie A. de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. Conservative plane releasing for spatial privacy protection in mixed reality. *CoRR*, abs/2004.08029, 2020.

[11] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

[12] Joseph L Gabbard, Gregory M Fitch, and Hyungil Kim. Behind the glass: Driver challenges and opportunities for ar automotive applications. *Proceedings of the IEEE*, 102(2):124–136, 2014.

[13] Sonja Grapemind. Smart glasses. https://thenounproject.com/term/smart-glasses/215218/. Accessed: 2021-05-21.

[14] Jassim Happa, Anthony Steed, and Mashhuda Glencross. Privacy-certification standards for extended-reality devices and services. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 397–398. IEEE, 2021.

[15] Scott Hayden. 'Rec Room' Studio Estimates Facebook Sold 2-3M Oculus Quest 2 Headsets in Q4. https://www.roadtovr.com/facebook-rec-room-2-3m-oculus-quest-2-sold/, 2021. Accessed: 2021-05-26.

[16] Brittan Heller. Watching Androids Dream of Electric Sheep: Immersive Technology, Biometric Psychography, and the Law. *Vanderbilt Journal of Entertainment & Technology Law*, 23(1), December 2020.

[17] Suman Jana, Arvind Narayanan, and Vitaly Shmatikov. A scanner darkly: Protecting user privacy from perceptual applications. In *2013 IEEE symposium on security and privacy*, pages 349–363. IEEE, 2013.

[18] Brendan John, Sophie Jörg, Sanjeev Koppal, and Eakta Jain. The security-utility trade-off for iris authentication and eye animation for social virtual avatars. *IEEE transactions on visualization and computer graphics*, 26(5):1880–1890, 2020.

[19] Brendan John, Sanjeev Koppal, and Eakta Jain. Eyeveil: degrading iris authentication in eye tracking headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–5, 2019.

[20] Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4309–4314, 2016.

[21] Kristen Lee. BMW wants customers to pay a subscription fee to use features the car already has installed, like a heated steering wheel or adaptive cruise control. https://www.businessinsider.com/bmw-subscription-model-for-features-2020-7, 2020. Accessed: 2021-05-27.

[22] 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)*, 2021.

[23] 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.

[24] Divine Maloney, Guo Freeman, and Andrew Robb. Social virtual reality: Ethical considerations and future directions for an emerging research space. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 271–277. IEEE, 2021.

[25] Divine Maloney, Samaneh Zamanifard, and Guo Freeman. Anonymity vs. familiarity: Self-disclosure and privacy in social virtual reality. In *26th ACM Symposium on Virtual Reality Software and Technology*, pages 1–9, 2020.

[26] Viraaji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.

[27] Arvind Narayanan and Vitaly Shmatikov. Robust Deanonymization of Large Sparse Datasets. In *IEEE Symposium on Security and Privacy*, 2008.

[28] Francesco Pittaluga, Sanjeev J Koppal, Sing Bing Kang, and Sudipta N Sinha. Revealing scenes by inverting structure from motion reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 145–154, 2019.

[29] Gerulf Rieger, Brian M Cash, Sarah M Merrill, James Jones-Rounds, Sanjay Muralidharan Dharmavaram, and Ritch C Savin-Williams. Sexual arousal: The correspondence of eyes and genitals. *Biological Psychology*, 104:56–64, 2015.

[30] Franziska Roesner, Tadayoshi Kohno, and David Molnar. Security and privacy for augmented reality systems. *Communications of the ACM*, 57(4):88–96, 2014.

[31] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J Wang. World-driven access control for continuous sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1169–1181, 2014.

[32] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. Secure multi-user content sharing for augmented reality applications. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 141–158, 2019.

[33] Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1310–1321, 2015.

[34] Piotr Siedlecki. Kid running. https://www.publicdomainpictures.net/en/view-image.php?image=126894&picture=kid-running. Accessed: 2021-05-21.

[35] 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.

[36] SunKing2. Desk - plain - perspective view. https://openclipart.org/detail/294639/desk-plain-perspective-view. Accessed: 2021-05-21.

[37] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.