

Augmenting Security and Privacy in the Virtual Realm: An Analysis of Extended Reality Devices

Derin Cayir, *Florida International University, Miami, FL, USA*

Abbas Acar, *Florida International University, Miami, FL, USA*

Riccardo Lazzeretti, *Sapienza University of Rome, Rome, Italy*

Marco Angelini, *Sapienza University of Rome, Rome, Italy*

Mauro Conti, *University of Padua, Padua, Italy*

Selcuk Uluagac, *Florida International University, Miami, FL, USA*

Abstract—In this work, we present a device-centric analysis of security and privacy attacks and defenses on Extended Reality (XR) devices, highlighting the need for robust, and privacy-aware security mechanisms. Based on our analysis, we present future research directions and propose design considerations to help ensure security and privacy of XR devices.

Extended Reality (XR) technologies stand at the forefront of a new digital revolution in an era marked by constant technological innovations. Nowadays, XR technology is much more than a device that produces 3D visuals. With new devices released each year and additional manufacturers getting involved in this field, the XR devices are considered for different application domains from entertainment to education to healthcare. The emerging metaverse realm offers a bright future with capabilities ranging from assisting astronauts in their mission to making hearing-impaired individuals "see" the conversations via subtitles.

XR devices are versatile in their functionality, equipped with an array of advanced sensors, communication capabilities, and hardware specifications. As these technologies evolve, our perception of reality seamlessly blends with the virtual world. However, the exponential growth of these technologies raises concerns about whether these devices are secure and the users' sensitive information is kept private. The increasing number of users will naturally attract attackers who will attempt to exploit these devices. The challenge arises from the diverse sectors currently utilizing these technologies and the unique properties

of the devices themselves. This heterogeneity of the devices aggregates the potential attacks, and complicates the examination of current devices. Thus, it is vital for the research community in this field and the developers of these devices to consider what the current technologies propose and the vulnerabilities that the attackers can exploit.

In this article, we study possible attacks on XR devices that could compromise the security and privacy of users and their environment in a device-centric approach. We highlight our key findings from detailed literature analysis, discuss the current attack vectors of XR devices, and present the security and privacy attacks with their corresponding defenses proposed in the literature. We analyze the attacks performed on the Virtual Environments (VE) separately, emphasizing the need for a further focus on this topic. Finally, we point to new research opportunities and propose design considerations, which can serve as valuable guidance for developers and the metaverse community.

Methodology

Literature Review: In order to find as many papers that perform security and privacy attacks or defenses on XR devices, we queried Google Scholar, ACM, and IEEE libraries on February 1, 2023. From 319 papers, we have restricted our selection to 41 papers listed in the table in Figure 4, testing practical attacks and




FIGURE 1. Spectrum of Extended Reality Technologies.

their defenses that target XR devices' security and privacy. For interested readers, we detail our literature review methodology and the PRISMA 2020 guidelines followed further in the GitHub repository³.

Device Search: From the selected papers, we gathered the devices used for the experiments. We also added to our device dataset other XR devices from the same companies that produced the devices mentioned in the papers. At the end of this process, in total, we identified 30 XR devices. The full list can also be found in the GitHub repository³.

Security and Privacy Analysis: We examined the devices' security and privacy by analyzing their documentation websites, manufacturer posts, articles, and blogs the links of which are given in the GitHub repository³. In addition to the on-device properties, we analyzed the literature to find information about the security and privacy vulnerabilities of the devices and what types of attacks were seen on them. The questions we discuss in privacy policies are "What type of data is collected, and where is it stored?", "Why is this data collected?", "Who is the data shared with?", "What are the users' rights on their data?" and "What are the privacy requirements of the apps on the devices?".

Security and Privacy Mechanisms in XR Devices

In this section, we examine some general properties of XR devices. Then, we highlight XR devices' security and privacy mechanisms using their security documentation and privacy policies.

General Properties of XR Devices

Virtual Reality (VR) aims to replace the real world with a digital world, fully separating the user from their surroundings. On the other hand, Augmented Reality (AR) overlays virtual objects onto physical objects in

the real world. Mixed Reality (MR) combines AR and VR, allowing interactive integration between the two worlds. XR encompasses AR, VR, and MR, containing all the devices that merge the virtual and real world, as shown in Figure 1. To seamlessly integrate the virtual world with the real world, XR devices strive to stimulate as many senses as possible (vision, hearing, smell, touch, and taste) through their sensors and actuators. Some general properties of XR that enhance realism are as follows:

Positional-tracking Features. XR devices offer 6 Degrees of Freedom (DoF) or 3 DoF, inside-out, or non-positional tracking. With 3 DoF the device can only track the rotational movement of the user, whereas with 6 DoF it tracks the user's rotation and position. These are achieved by the inbuilt sensors such as gyroscope, magnetometer, accelerometer, cameras, infrared sensors, and IMUs.

Tracking Sensors. Tracking sensors play a crucial role in XR devices and they span from tracking the users' motions and interactions to their environment.

User Motion Tracking: The XR devices have a Head-Mounted Display (HMD) that contains accelerometer, gyroscope, and magnetometer sensors to understand the head movements of the users. Similarly, the hand controllers are equipped with these motion sensors to track the position and orientation of the users' hands or even their finger movements. Many devices on the market, such as Meta Quest 2 and Microsoft HoloLens, support hand tracking where users can use their hands instead of a cursor. This is possible with inside-out cameras on the headsets⁴. XR devices can also detect the users' body motions, tracking different parts of the body to translate these movements into avatars. For example, HTC sells Vive Trackers, external devices that users can attach to their bodies to integrate their

movements into VR with more precise accuracy⁵.

User Interaction Tracking: Alongside translating body movements into the virtual realm, XR devices are equipped with eye and speech-tracking technologies which could be used to enhance the avatars, fully mimicking the users' speech and eye movements, and also developing more realistic simulations for medicine, missions, and much more. Microsoft HoloLens, Meta Quest Pro, and HTC Vive devices have eye-tracking sensors on the HMDs. With Vive Focus's eye tracker, users only need to gaze in a certain direction to open/close tabs or select objects⁶. Meta Quest Pro also captures and stores the raw face-image of the users to extract the user's natural facial expressions to create more natural-looking avatars.⁷

Environmental Tracking: XR devices have outward-facing cameras that track everything within the users' environment, and facilitate the precise rendering of 3D objects into the user's environments. Proximity sensors detect the presence of objects, while depth sensors enable the devices to create a 3D map of the users' environment. Although VR devices are not primarily designed to integrate real-world and virtual-world objects as AR/MR devices are, many contemporary VR devices, including Meta Quest 2, Pico 4, PSVR 2, and Magic Leap, still incorporate these sensors and pass-through cameras to enable room-scale inside-out tracking. Meta apps can also use pass-through cameras to blend the physical and virtual environment of the users, a purpose that goes beyond merely viewing and not processing the real environment's data⁸.

Audio and Speakers. Audio/speakers are integrated into the devices, and some devices have 3D spatial audio so users can physically locate the sounds they are experiencing in their virtual world. Meta Quest 2 and HTC Vive are examples of devices that use 3D spatial audio.

Haptic Feedback. Haptic feedback is an essential part of the VR experience to incorporate the senses of the users into their virtual world. There are different SDKs that support haptics for developing immersive apps such as vibrating the controllers⁹ and applying force to simulate touch. There are also additionally sold suits, and gloves designed to make the metaverse experience even more realistic.

Communications. The XR devices include WiFi and Bluetooth communication so that users can collaborate with other users or connect to their other gadgets. Each device has its own compatibility requirement and can run on different OS. For app development platforms,

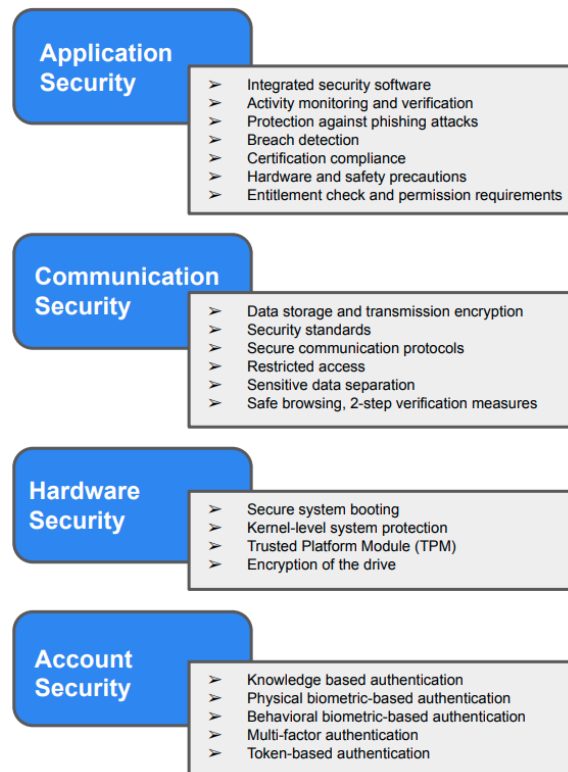


FIGURE 2. Security properties of XR devices.

the devices are compatible with different graphic cards and RAMs.

Security Properties of XR Devices

The impact of security and privacy attacks is high on XR devices as they are complex technologies that collect potentially user-identifiable information. Due to the immersive nature of these devices, attacks can manipulate users' perception of reality, potentially leading to physical harm. To ensure the security and privacy of the devices, vendors apply different methods that aim to meet the challenges of the modern cyber threats landscape, summarized in Figure 2.

Application Security. Applications are essential for delivering different functionalities to users. Since applications have access to users' sensitive data, securing them against exploitation of sensitive information is a high priority. Device vendors adopt various measures to achieve this goal. Microsoft HoloLens relies on Microsoft Defender SmartScreen, integrated into the OS, warning users of dangerous websites and applications that can perform phishing and malware delivery. Meta monitors and verifies the account activity to prevent

malicious acts and policy violations. Vuzix safeguard users' information from phishing attacks by preventing third-party apps from asking for user's sensitive information. Sony uses the information collected from the user to detect breaches such as unauthorized access to the apps. Pico Neo 4 uses "ETSI EN 303 645"-based security certification, which includes regular security updates and key management practices.

Communication Security. Device vendors use different encryption standards to prevent attackers from accessing sensitive user information. For instance, HTC has data processing or altering, anonymization, pseudonymization, encryption during transmission using TLS, and access restriction. However, as stated in their policy, HTC does not take responsibility for threats from independent third-party applications. Microsoft HoloLens 2 secures data transfer between itself and the cloud using Azure integration. Furthermore, Dynamics 365 Remote Assist helps when deploying to external clients, separating sensitive device vendor data and resources. Google devices ensure continuous encryption to keep data private while in transit and have security features like Safe Browsing, Security Checkups, 2-step verification, physical security measures, and restricted access to personal info. Quest's VR messenger app prioritizes security, testing end-to-end encryption and specific apps' access control with its latest updates. Similarly, Meta devices have end-to-end encryption. For digital audio and video content encryption, Epson Moverio supports HDCP-encrypted content. Additionally, the data in transit is protected through TLS by many other manufacturers, such as Magic Leap, Pico, Samsung, and Vuzix.

Hardware Security. Devices employ various hardware security measures to guard against unauthorized access and physical attacks. For instance, Microsoft HoloLens2 uses a Trusted Platform Module (TPM), a hardware-level security technology to generate, store cryptographic keys, and authenticate the device using unique RSA keys. Furthermore, BitLocker provides another level of security by encrypting the drive, employing AES-XTS-256 encryption, and safeguarding the data with multi-factor authentication, including Read Only (RO) media and privacy protection of writable data. Similarly, Pico devices use secure system booting, kernel-level system protection, and the implementation of a trusted environment.

Account Security. We note different authentication methods deployed on different devices. For instance, HTC Vive Cosmos and Epson Moverio use pin/password authentication. Quest Pro and Meta devices, on the other hand, use unlock patterns, also providing users with privacy customization options.

Furthermore, some devices use accounts for login. Such as, Pimax links its authentication to Steam account, Samsung authenticates through Oculus login, and Magic Leap relies on its own ID system where a code verification is sent to a registered email address. Similarly, PSVR devices use a QR code for the initial device sign-in and then four-digit passwords with two-step verification for the remaining entries. Microsoft HoloLens supports iris-based authentication, but the users can also choose password entry to log in to their devices. Some third-party apps utilize biometrics, such as PalmID, which stores encrypted biometric signatures in Epson Moverio devices. On the other hand, the Pico Neo 3 Pro only requires login for the Pico App Store due to its business-focused purpose, where setup and access to files and apps must be quick.

Privacy Policies of XR Devices

The built-in sensors in XR devices collect data during or after the usage of the gadgets. Many of today's devices collect and share this information according to their privacy policies. So, we discuss the privacy of the devices in the current market by examining their policies and summarizing their properties in this section.

What type of data is collected and how? Data collected by XR devices is highly sensitive, including information about users' physical properties, movements, environment, gender, age, gestures, and biometric information. If an attacker targets this data, the consequences can be damaging. Hence, users must be aware of the type of data the device vendors collect and where and how these data are stored. As stated in the Privacy Policies of Meta and HTC, the devices collect data in three ways: user-provided, sensor-collected, and third-party obtained. Information the users give while using devices may be about their transactions, social interactions, communications, email addresses, phone numbers, gender, location, physical features, avatar, content, and social media accounts. Automatically collected data may be about people, games, apps, and features the users interact with. Through cookies, the data is linked to the user, including information about product access, device type, IP address, unique identifiers, WiFi network, web traffic, environment, physical dimensions (e.g., height, head size), play area, hand size, and movement. Information gathered from third parties may be from apps, developers, content providers, and marketing partners.

Where is the data stored? Data storage practices vary between devices. Meta stores the data in the

device in its raw form. Similarly, Magic Leap 2 has no cloud nor centralized server connection and stores the data on the device. HTC stores the data on the user's phone or HTC's servers (encrypting the data and not transmitting it anywhere other than the device and connected PC).

Why is this data collected? Device vendors collect data for many purposes, including improving user experience, providing better-personalized services, communicating with the user, and protecting the manufacturers, its users, and the public (e.g., analyzing data to detect abuse, such as spam or illegal content). The data may also be used to enhance realism, such as using controllers, HMD movements, and audio to make the avatar more realistic.

Who is the data shared with? Data collected by the devices can be shared without the users' knowledge. It is crucial for the users to understand what is done with their data and for developers of these devices to know how other vendors handle the data they collect. Generally, the data is shared with domain administrators, advertisement network providers, affiliated companies, other users, and third parties with the users' consent. Many vendors state that the data may be transferred, stored, and processed in any other country the device manufacturers business in less protective privacy laws.

What rights do users have over their data? Users have the right to manage, update, limit, and delete their data as well as to oppose and withdraw consent for data collection and marketing messages as stated in Pico, HTC, and Vuzix's privacy policies. The user can do this by contacting the email provided on the website. Deleting a Meta account results in the deletion of posts, entities, and apps, but not other users' posts about that user. With PlayStation VR (PSVR) devices, users can adjust the amount of shared data through the settings.

What are the privacy requirements of the apps on the devices? Most devices analyzed in this paper are programmable, where at-home users can create their apps for their needs. However, this freedom comes with the cost of compromising the security and privacy of the devices. Developers should set basic app requirements to ensure a coherent experience and prevent malicious apps. Meta suggests Virtual Reality Check (VRC) guidelines for app developers in their Privacy Policy and requires the apps to follow their own privacy policies, linking to the policy and clearly explaining collected data and use. Similarly, Google proposes general rules that app developers must follow for the users' safety. They define what can be collected from the users and how the apps should form their own privacy and content policies. Moverio prohibits collecting any information without users' consent and

any phishing to gain sensitive information about the user.

Security Attacks and Defenses

In this section, we categorize the security attacks into two categories: 1) Attacks on XR Devices and 2) attacks via XR devices. Our attack categorizations are shown in Figure 3. Papers presenting the attacks are listed on Figure 4.

Attacks on XR Devices

Malware Attacks In this, an attacker plants viruses, or worms on users' devices without their knowledge. An example of a malware attack observed on VR devices is Big Brother, proposed by ReasonLabs.¹⁰ The malware can infect VR devices with Android-based OS. With this, the attacker can remotely connect to an Android-based VR device and record the headset screen. This malware infects the user's computer, and once the malware enters the PC, it waits for a Developer-Mode-enabled VR device to connect. Upon connection, it opens a TCP port to record the user's headset whenever the PC and VR device share the same WiFi network.

Also, ransomware can target XR devices, limiting users' access until a ransom amount is paid¹¹. An Android ransomware sample was tested on Meta Quest 2 by integrating Simple Ransomware Sample (SRS) on the device which is developed as a standard Android application [P1]. The goal was to get read and write data permissions through SRS, and encrypt the data with a function that uses Java Crypto and Security libraries. The researchers concluded that the attack surface of Meta Quest 2 includes essential elements that can be leveraged for effectively carrying out ransomware attacks.

Network Attacks In network attacks, the attackers exploit the vulnerabilities of the target network and bypass the security mechanisms in place. For instance, Valluripally et al. [P2] showed a Denial of Service Attack (DoS) was executed via packet tampering, duplication, and dropping, resulting in the crashing of virtual reality environment's server. Another study [P3] also showed that DoS attacks resulting in the frame-rate drops in the devices may lead to nausea and dizziness, and create cybersickness attacks.

Password Stealing Attacks Password-stealing attacks target the authentication of the devices and can lead to unauthorized access and sensitive information leak-

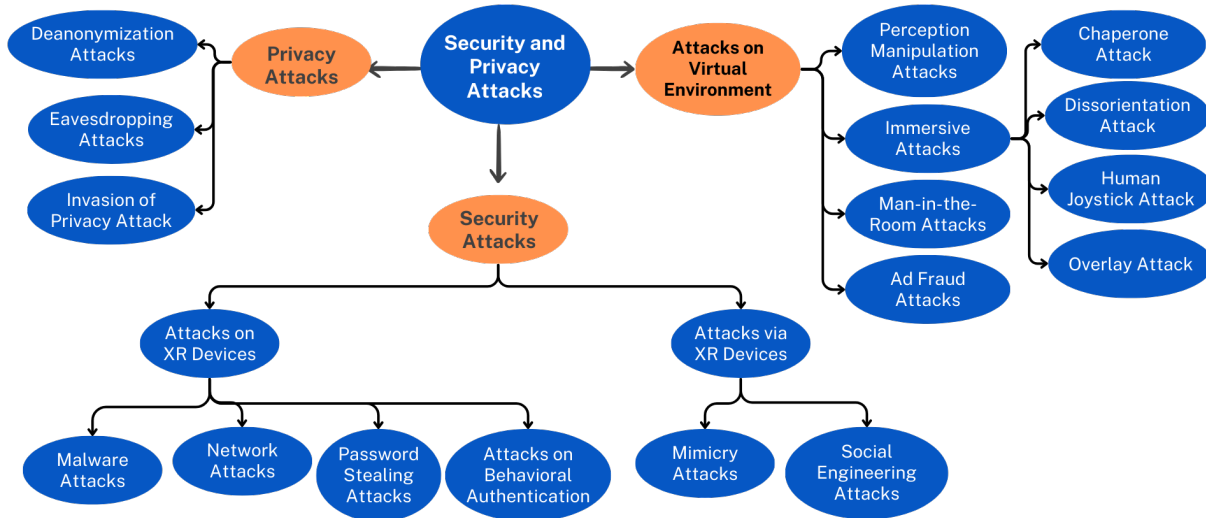


FIGURE 3. Security and privacy attacks in the literature.

age. Key-logging attacks can be performed by capturing users' hand traces to identify their passwords while using an in-air tapping keyboard for input [P4]. An adversary could plant hand tracker devices or videotape the users' text entry processes to obtain the victim's hand trace patterns and reconstruct inputs like passwords. This was also evident in a recent study where researchers retrieved graphical pattern lock inputs, passwords, emails, and pin entries of the users with VR HMDs, all from a video of the user interacting with the XR device [P5].

Moving beyond visual observations, other studies explored non-visual approaches to identify key-logging [P6]. These methods utilize a range of techniques, from analyzing network signals to leveraging device sensors, further highlighting the significance of this threat. For instance, Ling et al. [P7] performed vision-based and motion-based side-channel attacks on Samsung Gear VR devices using sensors. The motion-based side-channel attack, in particular, utilized the Samsung Gear VR's user motion tracking sensors by tricking the user into downloading a malicious app, which collects the orientation angles, hence giving information about where a key click occurs, and leading to the leakage of the user's password. Similarly, HMD's motion sensors from any XR devices with virtual keyboards were found to reveal characteristics for users' typing behavior, enabling to segment the motion signals and determine the typed words [P8]. Moreover, a recent paper further exploited side-channel information such as thread times to differentiate digit inputs using a spy program on Microsoft HoloLens 2 and Meta Quest 2 devices [P9].

Attacks on Behavioral Authentication Due to usability considerations, behavioral authentication systems are considered ideal for AR/VR devices. Miller et al. [P10] analyzed the ball-throwing task for authentication, where they could extract the 2D motion trajectories from the captured videos and match them to the 3D enrollment trajectories of users using HTC Vive, Vive Cosmos, and Meta Quest devices. Thus, they demonstrated that behavior-based authentication approaches could also be susceptible to attacks by obtaining the 2D video of the users.

Attacks via XR Devices

Mimicry Attacks AR devices can facilitate successful mimicry attacks on keystroke dynamics-like behavioral biometrics. Khan et al. [P11] proposed an AR-based approach to mimic touch dynamics used for smartphone authentication. The results showed that 87% of the attacks can bypass the authentication method.

Social Engineering Attacks The extensive data collection from advanced sensors of XR devices also poses security concerns through impersonation and social engineering attacks such as deep fakes. With deep fakes, an attacker could trick people into believing they are someone else. For instance, it has been shown that by creating deep fakes where the face and body of the user are physically altered to a digital form, the attackers can mimic the user and how they appear [P12]. To prevent deep-fake audios, [P13] suggests that developers can make a layer to modulate the voice input obtained from the HMD so that user's personal information is not identifiable.

#	Citation
P1	M. E. Mahan, "Exploring ransomware on the oculus quest 2," Ph.D. dissertation, Louisiana Tech University, 2022.
P2	S. Valluripally et al., "Modeling and defense of social virtual reality attacks inducing cybersickness," <i>IEEE Trans. on Dep. and Secure Comput.</i> , vol. 19, pp. 4127-4144, Oct. 2021.
P3	J. Happa et al., "Cyber security threats and challenges in collaborative mixed-reality." <i>Frontiers in ICT</i> 6, Apr. 2019.
P4	Ö. Meteriz-Yildiran et al. "A keylogging inference attack on air-tapping keyboards in virtual environments." in <i>IEEE Conf. on VR and 3D User Interf.</i> , 2022, pp. 765-774.
P5	S. R. K. Gopal et al., "Hidden reality: caution, your hand gesture inputs in the immersive virtual world are visible to all!" presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P6	A. Arafat et al., "Vr-spy: A side-channel attack on virtual key-logging in vr headsets." in <i>IEEE Conf. on VR and 3D User Interf.</i> , 2021, pp.564-572.
P7	Z. Ling et al., "I know what you enter on gear vr," presented at IEEE Conf. on Comm. and Network Sec., Jun. 10-12, 2019.
P8	C. Slocum et al. "Going through the motions:AR/VR keylogging from user head motions." presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P9	Y. Zhang et al., "It's all in your head (set): Side-channel attacks on ar/vr systems," presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P10	R. Miller et al., "Using external video to attack behavior-based security mechanisms in virtual reality (vr)," in <i>IEEE Conf. on VR and 3D User Interf.</i> , 2022, pp. 684-685.
P11	H. Khan, U. Hengartner, and D. Vogel, "Augmented reality-based mimicry attacks on behavior-based smartphone authentication," in <i>16th Ann. Int. Conf. on Mobile Sys.</i> , 2018, pp. 41-53.
P12	A. J. Bose and P. Aarabi, "Virtual fakes: Deepfakes for virtual reality," in <i>IEEE Workshop on Multi. Signal Proc.</i> , 2019, pp. 1-1.
P13	D. Maloney, S. Zamanifard, and G. Freeman, "Anonymity vs. familiarity: Self-disclosure and privacy in social virtual reality." in <i>ACM Symp. on VR Software and Tech.</i> , 2020, pp. 1-9.
P14	B. Falk et al., "Poster: reavatar: virtual reality de-anonymization attack through correlating movement signatures," in <i>ACM SIGSAC Conf. on Comp. and Comm. Sec.</i> , 2021, pp. 2405-2407.
P15	P. P. Tricomi et al., "You can't hide behind your head-set: User profiling in augmented and virtual reality," in <i>IEEE Access</i> , vol 11, pp. 9859-9875, 2022.
P16	V. Nair et al., "Unique identification of 50,000+ virtual reality users from head & hand motion data." presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P17	J. Li et al., "Kalcido: Real-Time privacy control for Eye-Tracking systems." presented at the 30th USENIX Sec. Symp., Aug 11-13, 2021.
P18	B. David-John et al. "Towards gaze-based prediction of the intent to interact in virtual reality." in <i>Proc ACM Sym. on Eye Track. Res. and App.</i> , 2021, pp. 1-7.
P19	B. David-John, et al. "A privacy-preserving approach to streaming eye-tracking data." <i>IEEE Trans. on Vis. and Comp. Graph.</i> , pp. 2555-2565, Mar. 2021.
P20	J. Steil et al. "Privacy-aware eye tracking using differential privacy." in <i>Proc. ACM Symp. on Eye Track. Res. and App.</i> , 2019, pp. 1-9.
P21	A. Liu et al. "Differential privacy for eye-tracking data." in <i>Proc. ACM Symp. on Eye Track. Res. and App.</i> , 2019, pp. 1-10.
P22	E. Bozkir et al. "Differential privacy for eye tracking with temporal correlations." <i>Plos one</i> 16.8, no. 8, 2021.
P23	Y. Kim et al. "Erebus: Access Control for Augmented Reality Systems." presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P24	C. Shi et al., "Face-mic: inferring live speech and speaker identity via subtle facial dynamics captured by ar/vr motion sensors," in <i>Proc. of 27th Ann. Int. Conf. on Mobile Comp. and Net.</i> , 2021, pp 479-490.
P25	R. Trimananda et al., "OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR." presented at the 31st USENIX Sec. Symp., Aug 10-12, 2022.
P26	T. Denning, Z. Dehlawi, and T. Kohno, "In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies." in <i>Proc. SIGCHI Conf. on Human Fac. in Comp. Sys.</i> , 2014, pp. 2377-2386.
P27	Lebeck, Kiron, et al. "Towards security and privacy for multi-user augmented reality: Foundations with end users." in <i>IEEE Symp. on Sec. and Privacy</i> , 2018, pp. 392-408.
P28	J. O'Hagan et al., "Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders" in <i>Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.</i> , 2023, pp. 1-35.
P29	H. Farrukh, et al. "LoCln: Inferring Semantic Location from Spatial Maps in Mixed Reality." presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P30	Y. Zhaoe et al., "Privacy-preserving Reflection Rendering for Augmented Reality.", in <i>Proc. ACM Int. Conf. on Multimedia</i> , 2022, pp. 2909-2918.
P31	J. De Guzman et al., "Security and privacy approaches in mixed reality: A literature survey." <i>ACM Computing Surveys</i> , vol. 52, no. 6, pp. 1-37, 2019.
P32	P. Casey et. al, "Immersive virtual reality attacks and the human joystick," <i>IEEE Trans. on Dep. and Secure Comp.</i> , vol. 18, no.2, 2019, pp. 550-562.
P33	F. Roesner, T. Kohno, and D. Molnar, "Security and privacy for augmented reality systems," <i>Commun. ACM</i> , vol. 57, no.4, pp. 88-96, 2014.
P34	K. Lebeck et al., "How to safely augment reality: Challenges and directions," in <i>Proc. 17th. Int. Workshop on Mobile Comput. Sys. and App.</i> , 2016, pp. 45-50.
P35	K. Ruth, T. Kohno, and F. Roesner, "Secure Multi-User content sharing for augmented reality applications." presented at the 28th USENIX Sec. Symp., Aug 14-16, 2019.
P36	S. Rajaram et al., "Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality." in <i>Proc. 2023 CHI Conf. on Human Factors in Comput. Sys.</i> , pp. 1-17.
P37	M. Vondrek, et al., "Rise of the metaverse's immersive virtual reality malware and the man-in-the-room attack & defenses," in <i>IEEE Comp. & Sec.</i> , 2022, p. 102923.
P38	K. Cheng et al., "Exploring user reactions and mental models towards perceptual manipulation attacks in mixed reality." presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
P39	H. Lee et al., "AdCube: WebVR Ad Fraud and Practical Confinement of Third-Party Ads." presented at the 30th USENIX Sec. Symp., Aug 11-13, 2021.
P40	Y. Abdrabou et al. "Understanding shoulder surfer behavior and attack patterns using virtual reality." in <i>Proc. 2022 Int. Conf. on Adv. Visual Int.</i> , 2022, pp. 1-9.
P41	F. Mathis, K. Vaniea, and M. Khamis. "Can i borrow your atm? using virtual reality for (simulated) in situ authentication research." in <i>IEEE Conf. on VR and 3D User Interf.</i> , 2022, pp. 301-310.

FIGURE 4. List of papers (additional references).

Privacy Attacks and Defenses

Privacy attacks focus on violating the users' right to exploit their personal information. In this section, we review the privacy attacks and defenses on XR devices.

De-anonymization Attacks

XR devices have many sensors that help navigate the user's environment, seamlessly blending real and virtual worlds. Naturally, these sensors collect information on the user's surroundings and personal information. Such information can be highly private as devices may record unique user movement data, potentially compromising their anonymity. For instance, in a de-anonymization attack called ReAvatar [P14], users are identified by their virtual avatar via correlating specifically recorded movements. Remarkably, the user's movement remains unique even when using multiple avatars so that attackers can also de-anonymize them across multiple avatars. Moreover, [P15] shows that AR (Microsoft HoloLens) and VR (HTC Vive Pro) platforms are vulnerable to de-anonymization attacks, by identifying the users from basic physical actions like walking and pointing. In a more recent study [P16], HMD and controllers' motion sensor data revealed patterns of user behavior, making potential attackers re-identify users across different sessions of popular games.

Another type of highly sensitive data is biometric data. Many XR devices - PSVR2, Magic Leap 2, Pimax Vision 8k, Microsoft HoloLens 2, and HTC Vive Pro Eye - widely adopt eye tracking technology for different purposes ranging from authentication to understanding users' interests for advertising. Given the rich information content eyes offer, this raises critical privacy concerns. For example, pupil size can be used to understand someone's interests, while eye movements can be analyzed to infer mental disorders, cognitive states, gender, and age [P17]. Researchers also found that the natural gaze dynamics from eye-tracking sensors could be used to predict the users' interaction with the virtual objects [P18], and also by attackers for user identification [P19].

Several strategies can be employed to safeguard user privacy, including the use of differential privacy, which involves the addition of random noise to obfuscate individual data without undermining overall data utility [P17], [P20-22]. Currently, independent content developers can directly access raw data collected by XR devices' sensors. Hence, to protect the users' privacy, the researchers proposed designing APIs that would add Gaussian noise to raw data while also implementing temporal and spatial downsampling [P19].

Furthermore, to prevent over-privileged malicious apps from accessing the raw sensor data, Kim et al. [P23] propose an access control scheme for AR which allows users limiting the access to sensor data.

Eavesdropping Attacks

An eavesdropping attack tested in HTC Vive Pro and Meta Quest devices is the Face-Mic approach that derives sensitive information by exploiting motion sensors [P24]. Speech-associated facial movements, bone-borne vibrations, and airborne vibrations of the user are collected, permitting to determine personal information, such as the user's gender. This attack utilizes zero-permission sensors (i.e., motion sensors) and reveals the user's protected information without the user's consent.

Invasion of Privacy Attacks

Invasion of privacy attacks involve the unauthorized collection and use of personal information. For instance, researchers realized that side-channel information could also be used for concurrent app fingerprinting on Microsoft HoloLens 2 devices [P9], identifying which app the user is currently using. Furthermore, several academic works found that the traffic flow from the XR devices revealed user-identifiable information, especially when the users were using social applications [P25].

Moreover, with outward-facing always-on cameras, the user themselves can record their environment in every detail without any notice, compromising the bystanders' privacy. Especially, AR glasses could be harder to notice in public settings, where bystanders might not expect to be recorded [P26]. Several papers conducted user studies to test bystander privacy experiences in crowded public spaces [P26-28], showing users' concerns about bystander privacy violations and invasive applications on their devices. In [P9], researchers found that environmental events created additional rendering, which was identifiable from performance counter analysis of the devices. From this, the researchers were able to identify the existence of a bystander by analyzing the CPU frame rates of Microsoft HoloLens, and they also were able to calculate the distance of the bystander from the device.

AR and MR devices capture spatial maps of the users' environment to overlay virtual content in the users' surroundings by depth sensors and always-on cameras, which introduce privacy concerns. Researchers found that this can reveal information about the location of the users [P29]. With a tailored malicious app, the researchers were able to extract the 3D

spatial map of the user's environment using Microsoft HoloLens and identified the user's indoor location from a model trained with 3D objects present in an environment. Another study [P30] found that inputs captured by AR devices during object rendering can contain sensitive objects, which will be translated onto the reflective AR objects. This reflection-based privacy attack results in the user's physical environment information being recovered by the attacker.

The literature suggests implementing defenses such as an intermediate layer between the sensor interfaces and the apps like input sanitization [P31]. This way, sensitive information can be protected by the input access control system. This can be achieved in two different ways:

Negotiating Permission: Developers can include an option where the bystanders have a right to opt-out if they feel their privacy is compromised [P31]. For instance, physical switches to block the cameras or push-pull notifications where the bystanders near an XR device receive an option not to get recorded can be implemented [P26].

Blurring: Developers can add a protection layer where sensitive objects (e.g., faces, license plates) in the captured images can be blurred [P26].

Attacks and Defenses in Virtual Environment (VE)

With XR devices, security and privacy concerns are not limited to the physical world. This section discusses the security and privacy issues in the VE.

Immersive Attacks

Immersive attacks target the unique properties of VR devices and are categorized into chaperone, disorientation, human joystick, and overlay attacks. A paper [P32] shows this is possible in HTC Vive and Oculus Rift devices by simply modifying VE parameters in a JSON file.

Chaperone Attack In a Chaperone attack, the attacker modifies the virtual boundaries of the victim [P32]. In situations where the user's confidence in the boundaries that are no longer valid is high, the attacker might do physical harm to the user by altering the boundaries. A proof of concept attack was performed, and HTC Vive and Oculus Rift devices were found to be vulnerable against all tested OpenVR and SteamVR applications [P32]. To perform the chaperone attack, the researchers obtained the artifacts such as the location of the VR boundaries, system settings,

and executable path location by exploiting SteamVR's vulnerability of storing the data in plain text without any integrity checks.

Disorientation Attack In the disorientation attack, the user's location and rotation were adjusted by making minor changes in the player's orientation through yaw and translation parameters [P32]. In cases where the users are immersed in virtual environments and subject to visual motion cues without physical motion, Visually Induced Motion Sicknesses (VIMS) are seen. This way, the player's orientation is controlled, forming a seasick sensation. Smaller fluctuations in the artifacts resulted in stronger seasick sensations. These attacks were performed through Steam and the success of this attack was similar to the Chaperone Attack as the same artifacts were targeted.

Human Joystick Attack Human Joystick attacks are designed to alter the direction or location of a user within the VE without their awareness [P32]. These attacks aim to manipulate the user's movement, potentially leading to physical harm, such as the user hitting an object. For instance, the virtual environment was shifted continuously to move the user to an attacker-defined location. To solve these attacks, some countermeasures are suggested: intrusion detection, where an attack is flagged if it detects any patterns different from the expected timing model, or securing timing information, where the modulation frequency of the optical signal is changed.

Overlay Attack Attackers can superimpose images (such as inappropriate or alarming content) onto the user's screen to potentially cause harm or distress or block the user's view. Loud songs can be played, and bright flashing lights can be displayed on the XR device, which may harm users physically. These attacks can be particularly dangerous because users may not realize that the overlaid content is not part of the XR experience and may react to it as if it were real. An unlimited number of images was overlaid on Oculus Rift and HTC Vive devices [P32]. Furthermore, it is also found that packet sniffing attacks can be used to capture the users' physical location parameters illegally to perform overlay attacks [P2].

Overlay attacks are also a valid concern for AR devices that are designed to overlay computer-generated visual, audio, and haptic signals onto the real world [P33]. In immersive AR applications, users must trust the app, and if it is targeted by the attacker, the users can be deceived about the real world. As a possible solution, windowing the display regions is

suggested where the OS gives the applications separate windows corresponding to the bounded regions of the display [P34]. With this solution, the applications' outputs are isolated from one another. Furthermore, [P34] proposes managing the outputs of AR devices as fine-grained objects, made of first-class OS primitives, which make the OS capable of controlling when and where objects are placed. This method yields better flexibility and output control than windowing the displays.

Security risks in AR do not just come from the apps themselves, but also from users who might intentionally spam others with disturbing virtual objects, or manipulate their virtual objects without permission. As a possible defense, Ruth et al. [P35] propose an app-level library or an OS interface tailored for AR multi-user application developers. They consider the users' expectations, who may have different expectations about how AR content should be shared. Their proposed framework sets security objectives for controlling other users' permission to access shared (outbound) content and managing the incoming (inbound) content and owned physical space. They introduce "ghost" objects where certain sensitive parts of the object are not shared with other AR users, and they suggest policies on physical space ownership in AR. Furthermore, Rajaram et al. [P36] pairs AR, Security and Privacy experts to find solutions to AR overlay attacks. This study highlights that virtual menus and proximity-based interactions were suggested for content sharing and access control techniques.

Man-in-the-Room Attacks

Man-in-the-Room (MITR) attacks represent a specific threat targeting the VE where users are known to share private information [P37]. These attacks often exploit the users' immersion within the VE, benefiting from their tendency to assume the same privacy norms that are valid in the real world also hold to the virtual world. For example, a private virtual room that users may use to communicate with each other may be targeted by a MITR attacker as users would feel secure in a virtual room and would not expect an outsider to join without their consent. However, via MITR attack, the attacker can exploit this perception and know everything happening inside a private VR room without the victim's knowledge or authorization [P37].

An example of MITR attacks was performed on the Bigscreen VR app on Steam which is supported by HTC Vive, Oculus Rift, and Windows Mixed Reality devices [P37]. The Bigscreen app is used for communication in a VR environment. The attackers

found a loophole where they exploited app vulnerabilities that caused a self-replicating infection (worm) without the user installing anything malicious. With the MITR attack, attackers could eavesdrop in the virtual room without other users noticing them. The attackers could turn on the users' microphones to listen to their conversations and observe what they were doing.

Perception Manipulation Attacks

Since the XR devices are designed to be highly immersive, many concerns have been expressed on the impacts of attacks on XR devices on the users. In [P38], the researchers created three attack scenarios targeting visual, auditory, and situational awareness perceptions.

With the visual attacks, the researchers overlay an adversarial virtual object, observing that the participants were fooled into believing the overlaid content was real, and their reaction times were significantly slowed. Interestingly, after the presence of the attacks, the participants started becoming hesitant and getting triggered by non-adversarial content. Imagine a real-world scenario where a user uses an XR device to get real-time guidance when driving and an adversary overlays incorrect speed limits, and traffic signs. The user will be deceived by these overlays and have a reduced reaction time, which is a valid concern while driving. The issue is that these attacks' impact continues even after the attack is finished as the users will lose their trust in the device and become hesitant with each traffic sign encountered.

Auditory attacks were performed while the users were concentrating on memorizing a sequence of elements. The immersive nature of XR audio led the users to treat the audio cues as a real-world stimulus. Lastly, the researchers displayed notifications on a screen which is in the background and realized that participants were not quick to notice real-world instructions while using their XR devices, showing that users are more focused in the metaverse than the real world.

Ad Fraud Attacks

Web-based VEs can be targeted by adversaries to create ad frauds that generate unintended ad traffic involving ad impressions or clicks [P39]. In XR technologies, the 3D world is rendered on an HTML canvas document object model (DOM) to create immersive experiences for the users and help them interact with the web page they are browsing. There are currently no primitives to separate the execution of an ad-serving JS script, enabling researchers to launch different attacks. One of these attacks was called gaze and con-

troller jacking attacks, where a fake gaze and controller cursor were created to make the users intentionally click on the malicious VR objects. Furthermore, with a blind spot tracking attack, the researchers exposed the limited visual awareness of the users during 360-degree views by placing malicious promotional objects in the blind spots of the users' views. Similarly, with the abuse of an auxiliary display attack, the researchers could block the user from seeing their immersive world by displaying ads. As a potential solution, the researchers propose AdCube, which sandboxes the ad-serving JS and suggests that ad entities should be given a confined area. The researchers also suggest publishers specify DOMs that interact with a confined third-party ad script and generate access control policies on write and read permissions for DOMs.

Future Research Directions

In this section, we leverage the insights gleaned from our study.

Authentication is the leading defense method. The current literature proposes unique ideas for user authentication, ranging from behavioral methods such as throwing a virtual ball to biometrics that utilize almost every part of the human body. Although authentication methods are the basis for securing the device from outsiders, because none of the devices have adopted the proposed authentication methods, it is clear that authentication offers only a partial panacea for device security.

Future Research Direction #1: Authentication alone cannot guarantee complete security, and it is important to consider multiple layers of security to address all possible attack vectors. Therefore, researchers must propose additional defense strategies that tackle a broader range of security threats and vulnerabilities.

XR devices as virtual testbeds. Alongside XR devices serving as tools for various security attacks, they can also be used to create realistic virtual testbeds. This idea is explored in academia by generating scenes in XR devices to understand attacker behaviors [P40] and test the proposed methods' usability [P41]. VR-generated test environments provide remarkable similarities to real-world scenarios while addressing the shortcomings of in-person studies, such as overcoming ethical and legal constraints. Given their inherent flexibility, VEs are easily modifiable, making them ideal for such testing and educational scenarios.

Future Research Direction #2: Professionals across diverse disciplines can utilize the XR devices to generate realistic testbeds and evaluate their algorithms within a remarkably authentic, yet controlled, environment. Additionally, VEs can facilitate testing the usability and efficiency studies of the defense solutions on the users.

Device diversity in security testing. The researchers predominantly utilize the same devices to apply their findings. The most used products were HTC Vive and Meta Quest 2 due to their wide accessibility and general public use. While we cannot assert that other devices not mentioned in this article are fully secure, we recommend the readers focus on OS characteristics or examine the root causes of the vulnerabilities when understanding whether a type of attack is also applicable to their devices.

Future Research Direction #3: Future research should conduct security assessments using several devices, beyond just the popular ones. This way, more attack vectors can be uncovered, identifying new potential vulnerabilities in a rapidly developing field.

Design Considerations

This section presents practical guidelines from our study to help developers create safer, more secure XR devices.

Protection of sensitive data. The immersive experience the XR technologies provide is made possible through the advanced sensors they are equipped with. However, our findings highlight that XR devices pose security and privacy risks by collecting intrusive sensor data, which can also expand the attack surface for other devices.

Design Consideration #1: The accessibility of raw sensor data within XR device app development environments has established a notable threat model. Therefore, we recommend developers of commonly used app development platforms (e.g., Unity, Unreal Engine) to incorporate a default setting that limits the accessibility of users' raw data to independent app developers. Such as, implementing differential privacy measures would protect the user data without compromising the app's performance.

Physical input methods. Input methods for sensitive data (e.g., passwords, text messages, emails) are highly physical as the user points the hands to a predefined location on a virtual keyboard. This opens up XR devices to numerous attacks, wherein an attacker will potentially extract the users' key presses, or replicate the authentication method by observing their actions.

Design Consideration #2: To prevent attackers from inferring the users' inputs, developers should utilize non-physical input methods. Eye-tracking technologies could be used for users to enter their passwords where they will enter their keys by looking at a key for a predetermined amount of time. Additionally, developers might consider methods like shuffling the keys of the keyboards to avoid virtual keyboard password-stealing attacks.

Virtual Environment as a new attack vector. Security and privacy issues such as MITR attacks or inferring user passwords through user motions are specific to targeting the VE of a user.⁵ While using XR devices, a user must continuously trust the environment generated by the devices. Hence, when an attacker targets the VE, the user who is fully immersed will be drastically affected.

Design Consideration #3: In the design phase of virtual environments of VR and MR devices, the developers and device manufacturers should incorporate user feedback mechanisms. Utilizing insights from user studies on VEs, such as the one conducted by Lebeck et al. [P27], can provide an understanding of the users' needs, and expectations from the VEs. Additionally, direct features like in-app surveys can be done to further enhance user security.

Vague privacy policies. Several vendors' privacy policies are not explicitly tailored to individual devices and fail to distinguish between the data collected when using an HMD and other scenarios. Moreover, in current privacy policies, there is no explicit identification of who among the partners, developers, domain administrators, or affiliated manufacturers the data is shared with.

Design Consideration #4: Sensitive data collection by XR devices requires clear communication and transparency from developers and manufacturers to users. Therefore, manufacturers should make their privacy policies easily accessible and understandable, communicating transparently about data collection and management processes. Features like opt-out options and data collection indicators should be added.

In this article, we focused on the emerging technology of XR, conducting a comprehensive analysis of the security and privacy mechanisms of the devices currently dominating the market. Specifically, we provided an evidence-based approach where we analyzed the literature for security/privacy attacks on XR devices. We have also highlighted the critical need to analyze attacks and defenses in the VE. Lastly, we provided the lessons learned, which discuss the topics that could be further explored as future research and suggest some design considerations for developers to improve the security and privacy of their applications. Overall, this paper aims to help researchers understand what is currently needed as future defense directions and take appropriate measures.

ACKNOWLEDGMENTS

The author(s) would like to thank A, B, and C. This work was supported by XYZ under Grant ###.

REFERENCES

1. *CSRankings: Computer Science Rankings*, CSRankings, Sep. 2023. [Online]. Available: <https://csrankings.org/#/index?all&us>
2. S. Stephenson et al., "SoK: Authentication in Augmented and Virtual Reality," in *Proc. IEEE Symp. on Sec. and Privacy*, 2022.
3. *Methodology of "Augmenting Security and Privacy in the Virtual Realm: An Analysis of Extended Reality Devices"*, Github, Sep. 2023. [Online]. Available: https://github.com/csfiu/Augmenting_Security_and_Privacy_in_the_Virtual_Realm-Methodology
4. *Hand Tracking Privacy Notice*, Meta, Sep. 2023. [Online]. Available: <https://www.meta.com/help/quest/articles/accounts/privacy-information-and-settings/hand-tracking-privacy-notice/>
5. *Introducing Vive Tracker*, Vive, Sep. 2023. [Online]. Available: <https://www.vive.com/us/accessory/tracker3/>
6. *Vive Focus 3 Eye Tracker*, Vive, Sep. 2023. [Online]. Available: <https://business.vive.com/eu/product/vive-focus-3-eye-tracker/>
7. *Natural Facial Expressions Privacy Notice*, Meta, Sep. 2023. [Online]. Available: <https://www.meta.com/help/quest/articles/accounts/privacy-information-and-settings/natural-facial-expressions-privacy-notice/>
8. *Use Passthrough on Meta Quest*, Meta, Sep. 2023. [Online]. Available: <https://www.meta.com/help/quest/articles/in-vr-experiences/oculus-features/passthrough/>
9. *Haptic Feedback*, Meta, Sep. 2023. [Online]. Available: <https://developer.oculus.com/documentation/unity/unity-haptics/>
10. *"Big Brother": A New Attack Vector Affecting Metaverse Security*, ReasonLabs Research Team, Sep. 2023. [Online]. Available: <https://reasonlabs.com/research/big-brother>
11. H. Oz et al., "RøB: Ransomware over Modern Web Browsers," presented at the 32nd USENIX Sec. Symp., Aug 9-11, 2023.
12. F. Roesner, T. Kohno, and D. Molnar, "Security and privacy for augmented reality systems," *Commun. ACM*, vol. 57, no.4, pp. 88-96, 2014.
13. K. Ruth, T. Kohno, and F. Roesner, "Secure Multi-User content sharing for augmented reality applications." presented at the 28th USENIX Sec. Symp., Aug 14-16, 2019.
14. S. Rajaram et al., "Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality." in *Proc. 2023 CHI Conf. on Human Factors in Comput. Sys.*, pp. 1-17.
15. T. Denning, Z. Dehlawi, and T. Kohno, "In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies." in *Proc. SIGCHI Conf. on Human Fac. in Comp. Sys.*, 2014, pp. 2377-2386.

Derin Cayir, is a Ph.D. student at Florida International University (FIU), USA. Her research interests include privacy/security systems for XR devices. She is working as a Graduate Research Assistant in Cyber-Physical Systems Security Lab (CSL). She is an IEEE CS Student Member. Contact her at dcayi001@fiu.com.

Abbas Acar, is a Postdoctoral Associate at the CSL at FIU, USA. His research interests include privacy-aware technologies, alternative authentication methods, and security/privacy issues related to IoT. Contact him at aacar001@fiu.edu.

Marco Angelini, is an Assistant Professor in Engineering in Computer Science at Sapienza University of Rome, Italy. His research interests include Visual Analytics, applied in the Cybersecurity domain. Contact him at angelini@diag.uniroma1.it.

Riccardo Lazzeretti, is an Associate Professor at Sapienza University of Rome, Italy. His research focuses on security and privacy, with a particular focus on IoT. He is an IEEE senior member. Contact him at lazzeretti@diag.uniroma1.it.

Mauro Conti, is a Professor at the University of Padua, Italy, and he is also affiliated with TU Delft and the University of Washington, Seattle. His main research interest is in the area of Security and Privacy where he published more than 500 papers in topmost international peer-reviewed journals and conferences. He is a Fellow of the IEEE, AAIA, and Young Academy of Europe, and a Senior Member of the ACM. Contact him at mauro.conti@unipd.it.

Selcuk Uluagac, is currently an Eminent Scholar Chaired Professor in the Knight Foundation School of Computing and Information Science at FIU, where he leads the CSL. His research focuses on cybersecurity and privacy with practical and applied aspects where he holds hundreds of research publications in the most reputable venues. Contact him at suluagac@fiu.edu.