

SoK: Data Privacy in Virtual Reality

Gonzalo Munilla Garrido*

TU Munich
Munich, Germany
gonzalo.munilla-garrido@tum.de

Vivek Nair

UC Berkeley
Berkeley, USA
vcn@berkeley.edu

Dawn Song

UC Berkeley
Berkeley, USA
dawnsong@berkeley.edu

ABSTRACT

The adoption of virtual reality (VR) technologies has rapidly gained momentum in recent years as companies around the world begin to position the so-called “metaverse” as the next major medium for accessing and interacting with the internet. While consumers have become accustomed to a degree of data harvesting on the web, the real-time nature of data sharing in the metaverse indicates that privacy concerns are likely to be even more prevalent in the new “Web 3.0.” Research into VR privacy has demonstrated that a plethora of sensitive personal information is observable by various would-be adversaries from just a few minutes of telemetry data. On the other hand, we have yet to see VR parallels for many privacy-preserving tools aimed at mitigating threats on conventional platforms. This paper aims to systematize knowledge on the landscape of VR privacy threats and countermeasures by proposing a comprehensive taxonomy of data attributes, protections, and adversaries based on the study of 74 collected publications. We complement our qualitative discussion with a statistical analysis of the risk associated with various data sources inherent to VR in consideration of the known attacks and defenses. By focusing on highlighting the clear outstanding opportunities, we hope to motivate and guide further research into this increasingly important field.

KEYWORDS

Augmented reality, security, threat model, defense model, attacks

1 INTRODUCTION

Virtual reality (VR) has recently become a major investing target for leading tech industry players aiming towards the so-called “metaverse” [118], a paradigm shift towards an internet in the form of a 3D virtual world. This new internet would require VR devices such as headsets and hand-held controllers to digitize and relay users’ physical characteristics and movements to other users worldwide for immersive interaction. While the “metaverse” might hold promise, researchers have recently shown how easily an attacker could violate data privacy by identifying [22, 90, 105] and profiling [92, 126, 142] VR users with a few minutes of data streaming, and demonstrated that the scope and scale of data collection in VR supersedes the capabilities of current internet platforms [92]. Researchers further illustrated how malicious developers could hide artifacts in virtual environments that inconspicuously induce

users to reveal personal information, e.g., by playing a seemingly innocent game [92]. These attacks are partly possible due to VR’s unparalleled immersiveness, which can make users more susceptible to self-disclosure [78, 135], and social engineering [3, 36].

Unlike the current internet platforms, where users can employ Tor, VPNs, proxies, and incognito mode to fend off user tracking and profiling, there is no equivalent and mature defense suite for combating unique threats in VR. The extant literature offers a scattered set of privacy defenses at a proof-of-concept stage with no significant knowledge transfer to commercial-grade applications. Industry practices are not encouraging either, as VR devices showed vulnerabilities [100], some developers ignore their own privacy policies, and their updates trend towards more data harvesting [143].

The looming unprecedented privacy threats in VR entail a “dystopian metaverse” as long as we lack the robust defenses available on the current internet. With this study, we aim to begin tackling this impending challenge by providing new, comprehensive threat and defense models, and classifications of data attributes, attacks, and defenses at the intersection of VR and privacy. Specifically, our SoK evaluates, systematizes, and contextualizes 75 studies filtered from over 1700 publications. The structure of the SoK flows naturally in the context of privacy in VR: Sensitive data attributes are targeted and harvested by adversaries. In response, privacy practitioners design the corresponding defenses. Each taxonomy we introduce builds upon the previous classification, e.g., we classified VR attacks based on our proposed threat model and attributes taxonomy. Lastly, we outline privacy opportunities by highlighting some of the most vulnerable yet easiest-to-protect attributes in VR.

With the following **contributions**, we hope to guide researchers to find the right privacy opportunities quickly:

- (1) We proposed a holistic information flow, threat, and defense models to frame future studies (§4).
- (2) We built a taxonomy of data attributes collectible in VR clustered in 9 classes with more than 60 data points (§5).
- (3) We categorized 35 attacks (§6) and 35 defenses (§7), and focused our insights by answering 10 research questions.
- (4) We performed a quantitative and qualitative analysis to outline practices and opportunities (§8), and extracted 12 findings and future work items to advance VR defenses (§9).

2 BACKGROUND & RELATED WORK

In the context of the reality and virtuality continuum of Milgram and Kishino [89], our SoK focuses on virtual reality (VR), i.e. a computer-simulated interactive environment experienced in the first person [121] (see Appendix A). The considered VR devices stem from reviewing the 75 collected publications.

*Corresponding author.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2024(1), 21–40
© 2024 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2024-0003>

2.1 VR Devices

Since the wave of mature VR products of 2016 (see Appendix B for the compiled list), the wider public has experienced immersive VR like never before. Primarily, users employ a head-mounted display (HMD) with (integrated) speakers, a microphone, and two handheld controllers with buttons [83]. Some HMDs tether to a PC [148], while others can remain wireless [83]. The VR system tracks the HMD and the controllers by outside-in tracking (using stationary external sensors [146]) or inside-out tracking (employing built-in optical sensors and inertial measurement units: three-axis accelerometers and gyroscopes [83]). Front cameras for inside-out tracking also enable the user to observe their real-world surroundings. This basic setup generates realistic 3D graphics, spatial audio, verbal interaction, and six degrees of head and hand tracking (X, Y, and Z positions, and yaw, pitch and roll) Manufacturers design VR devices for use in a “controlled” environment (e.g., home, backyard, office, etc.).

Other devices make VR experiences more immersive yet more pervasive, as they require additional sensitive user input. Optical sensors for eye-tracking enable foveated rendering [149], increasing the quality of the visual output [2] and lengthening HMD battery life by reducing GPU load. Moreover, eye-tracking combined with additional optical sensors that register facial features (face trackers [145]) enables telepresence via expressive (photorealistic) avatars [23]. Handheld controllers (and HMDs [112]) can provide haptic feedback, have touch sensors for detecting holding gestures [147], and the latest models have outward cameras for improved tracking [84]. A natural transition is forced feedback gloves that provide more ergonomic and realistic interactions [54]; in contrast, users can also employ conventional keyboards [124]. Full-body tracking [127] enables more expressive and richer experiences with other users in virtual worlds. Additionally, users may dawn haptic vests [15] that deliver positional haptic feedback and prototypical masks that emulate smells [107]. Healthcare VR applications include sensors that measure galvanic skin response [58], electrodermal activity [4], heart rate [137], skin temperature [106] and measure brain waves (HMDs with EEGs for brain-computer interfacing) [99]. This plethora of sensors and feedback devices facilitate immersive digital interactions in VR, but also pose significant privacy concerns due to the potential exposure of sensitive user data, such as biometrics, behavior, identity, and real-world surroundings [40, 90, 92, 143].

Augmented Reality (AR) and VR share several underlying principles but present distinct challenges in the context of privacy and security. Appendix D delineates the characteristics that set them apart in terms of security and privacy. In contrast, these privacy concerns diverge considerably from those found in smartphone and wearable ecosystems. While smartphone and wearable apps primarily utilize motion sensors, location data, and a few biometrics, VR captures a richer spectrum of user behaviors, movements, and emotions, resulting in a broader threat surface. Beyond the simple metrics like step counts on smartphones, VR systems interpret nuanced user behaviors and emotional responses. This level of immersion is not just an extension of smartphone privacy challenges; it embodies a new dimension of user privacy, largely due to VR’s capability to consolidate attributes previously scattered across multiple devices.

2.2 VR Attacks & Defenses

There are early seminal reviews, surveys, and SoKs on mixed reality (MR) displays (1994) [89], classifications (1996) [13], early challenges (1997) [8], integrity and ownership (2000) [46], and enabler technologies (2001) [9]. In the 2010s, researchers delved into ethical considerations of MR (2014, 2018) [1, 56], updated its challenges (2016) [115], discussed the threats of converging VR and social networks (2016) [101], and studied VR safety (e.g., epilepsy) (2018) [10]. Most recently, practitioners have continued work on VR safety attacks (e.g., misleading users to collide with their real-world surroundings) (2021, 2022) [21, 144], malicious VR ads and protections [69] (2021), user authentication (2022) [39, 134], and advocated for new regulations for the upcoming metaverse applications (2022) [121]. While the ethics, authentication, safety, ads, regulations, and underlying technologies are paramount for VR, these works lack a focus on privacy attacks and defenses.

Our systematic search revealed the most relevant related work at the intersection of VR and privacy. We collected 12 relevant literature reviews (LR) [20, 31, 35, 36, 40, 49, 50, 65, 67, 102, 128, 153], three of which are the closest to our work. Shrestha and Saxena (2017) [128] provided an offensive and defensive overview of eye-wears and HMDs with a focus on optical cameras in the fields of privacy, security (authentication and device integrity), and safety, with an emphasis in the latter two. De Guzman et al. (2019) [31] expanded the augmented reality (AR) privacy and security defense classification of Roesner et al. [120] to MR without an in-depth analysis of data attributes and attacks. Odeleye et al. (2022) [102] provided a taxonomy of cybersecurity VR attacks related to authentication and privacy, comprising 5 privacy defenses and 10 attack-focused studies, which we also included in §4, §6, and §7. In contrast, our work presents a more detailed exposition specific to VR and privacy, and, yet, contains a more comprehensive taxonomy of vulnerable data attributes, attacks, and defenses and a technical component on user data to highlight privacy opportunities (unlike any other LR). Overall, extant literature examines VR and privacy as a subset of broader reviews in MR, security, and safety, thus, our detailed study focused on VR and privacy is not part of prior work.

Among the rest of selected LRs, three delved in a specific sub-field of VR and privacy, specifically, Katsini et al. [65], Kröger et al. [67], and Gressel et al. [50] studied the privacy implications and research directions of eye-tracking. In our work, we compiled their findings in our comprehensive taxonomies. Additionally, we included the relevant privacy-related insights and VR application taxonomies of two comprehensive LRs that covered general metaverse topics as varied as data management, privacy, legal issues, and economic threats [40, 153]. Lastly, we included key information from narrower surveys in security and privacy in VR [35, 49], and data attributes and user privacy considerations [20, 36].

3 METHOD: DATA COLLECTION & ANALYSIS

The following summarizes our search approach and results described in detail in Appendix C. We employed seven of the most relevant digital libraries focused on computer science and software engineering in combination with Google Scholar to perform an exhaustive search of the extant literature. We only included publications containing taxonomies of VR data attributes or applications,

or aimed to review or implement privacy attacks or defenses in VR, from which we extracted the relevant artifacts to construct our comprehensive models and taxonomies.

With a curated set of keywords from our base literature of 12 publications, our initial systematic search generated 1700 hits, which we discussed and filtered (by title, abstract, and body), resulting in only 16 studies. Combining the base literature (12) with the filtered studies (16) resulted in 23 selected studies after deduplication. We then queried their authors for additional relevant work, performed a backward search of the references of the 23 studies, and added publications found thereafter, obtaining a final body of 75 publications—the most recent study dates May 2023. Further, to focus our VR attacks and defenses discussion, we designed 10 research questions (RQ).

Lastly, we complemented the taxonomies with an analysis of VR privacy opportunities (§8) and key findings and future work items (§9). Part of the analysis quantifies some of the most sensitive and easiest-to-protect data attributes by calculating a PCA of inference attacks and weighted mean accuracy degradation after enabling the corresponding defenses. We replicated these attacks and defenses from the most comprehensive frameworks among the 75 studies.

4 VR THREAT & DEFENSE MODELS

Method. From the 75 selected studies, we identified 5 studies that proposed a VR information flow [28, 40, 49, 92, 153] and 23 with an explicit discussion or proposal for VR defense and (predominantly) threat models [7, 27, 29–32, 40, 45, 59, 60, 62, 70, 72, 74, 92, 105, 125, 126, 130, 131, 139, 150, 161]. Two researchers extracted, discussed, and combined the associated artifacts, resulting in a holistic VR information flow that frames our threat and defense models.

4.1 VR Information Flow

VR device manufacturers or vendors provide app stores where users can download VR applications and games (e.g., from the Oculus Store or SteamVR). Fig. 1 illustrates the information flow after installing such an application. These applications typically run in the host VR system, which ingests user input: *geospatial & inertial data, audio, text, video, and physiological signals* (1A). The various VR devices process raw sensor data and other input types into useful telemetry, which the application accesses via an API (e.g., OpenVR) (2A). Such application controls how to use this data to generate different stimuli, e.g., visuals via a graphics rendering pipeline, audio through speakers, and haptics using hardware such as feedback vests (2B). The output devices present this processed information to the user as an immersive, interactive virtual world (1B). For multi-user online experiences, the client-side application exchanges processed telemetry with an external server through a network, which can reveal *system* and *network* user-specific information (3). Finally, the server updates the global state of the virtual world and relays telemetry to other users (4). As the information flows from steps (1A) to (4), intermediate data processing steps like filtering and compression degrade data quality.

4.2 VR Threats

Within the frame of this study, we consider a state of *privacy* as the lack of a breach of any individual’s sensitive data attributes [158].

In our threat model, attackers breach user privacy by collecting and inferring enough information to reliably *identify* and comprehensively *profile* a user across VR applications over multiple usage sessions (tracking). Attackers (i) *identify* an individual when they can uniquely distinguish the user from others, and (ii) *profile* users when they unwarrantedly attach information related to the user’s characteristics (e.g., demographics, preferences, etc.) [34, 64, 142].

The collected studies discussing or proposing threat models consider application developers [27, 29, 30, 32, 45, 59, 60, 70, 72, 74, 92, 126, 130, 139, 161], servers [7, 30, 92, 123], content creators [40, 92], device manufacturers [92, 131], other users [69, 105, 125, 150], and hackers¹ [40, 62, 139] as the attackers in VR, or rely on general privacy threat models like Lindunn [31, 34, 64, 70]. Based on these studies and their system decomposition, we adopt a more comprehensive and pervasive privacy-centered attacker classification for VR that encompasses the privacy repercussions of the above threat models. The adversary types of Fig. 1 correspond to four distinct entities associated with data processing in VR applications at different privilege levels. These adversaries might coalesce, e.g., a developer of a VR application can also run the server providing multi-user functionality. Table 1 shows these attackers’ capabilities.

While hardware and client attacks are not exclusive to VR, the unique capability of VR devices to amalgamate diverse data points, which traditionally necessitated multiple devices, underscores their distinct relevance in this context. Lastly, note that there is a spectrum of adversarial capabilities across adversary types and within each adversary category, whose strength in practice might be lower than our strong adversary model, e.g., by permission walls.

(I) Hardware Adversaries control the hardware and firmware of the VR device and, thus, can access raw user inputs and arbitrarily manipulate the information provided to the application (2A) and presented to the user (1B).

(II) Client Adversaries represent the developers of the client-side VR application running on the VR device (*Application Adversary* [92]) and the content creators (*Content Adversary* [3, 69]). Content adversaries can create *immersive falsehoods*, i.e., designing immersive experiences with misinformative, manipulative, and deceptive content [3, 69]. Application adversaries can access the input data via system APIs, and arbitrarily manipulate the rendered frames and signals output to the VR devices (2B) and the information streamed to external servers (3). Further, the current push towards VR multitasking [86], where different apps are running concurrently, sharing the output or competing for user’s inputs, increases the attack surface of the client adversary.

(III) Server Adversaries oversee the server enabling multi-user functionality and can arbitrarily process networked data before streaming it to users (4). This adversary encompasses the role of a potential *network* adversary by intercepting network traffic during steps (3) and (4) of the information flow [7].

(IV) User Adversaries represent other users of the same VR application. They receive user data streams from a server and can

¹Hackers can attack VR devices, servers, networks, databases, or perform shoulder-surfing—covered extensively by security literature [52, 108, 109, 109, 119]. These *security attacks* are depicted in Fig. 1.

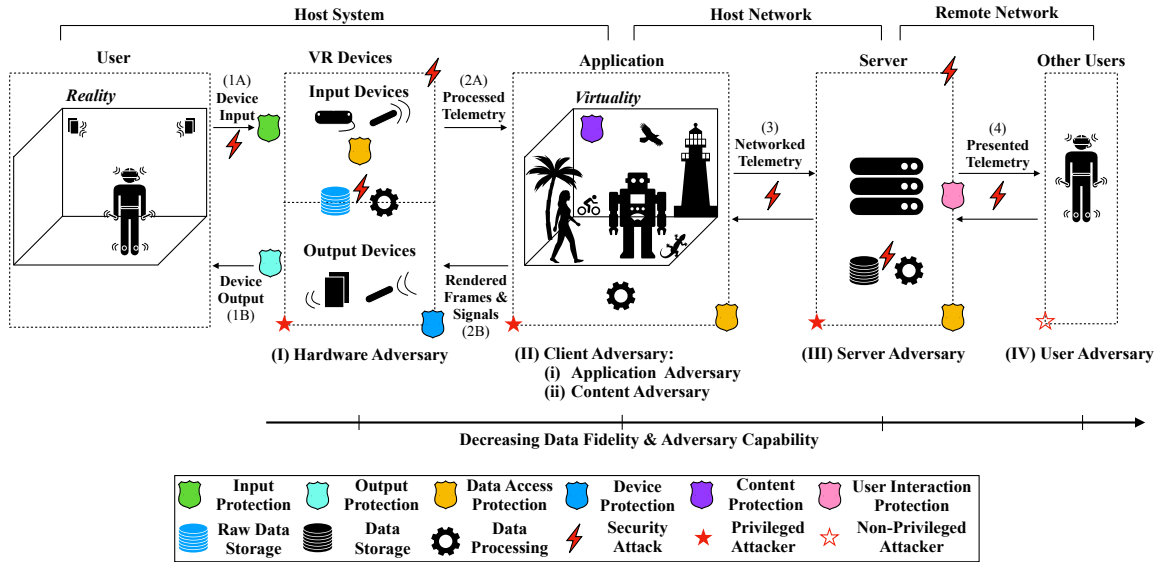


Figure 1: VR information flow depicting our defense model protections, adversary model attackers, and cyber attack points. (cf. [28, 92, 93, 153]).

interact with the target user. This attacker absorbs a potential *external observer* adversary executing shoulder surfing or other automated observational attacks during step (1A) of the information flow [74, 132].

Table 1: VR attacker capabilities (cf. [92]).

Adversary	Observable Attribute Classes							
	Geo. & Iner. Telemetry	Text	Audio	Video	Phy. Signals	System	Network	Behavior
(I) Hardware	✓	✓	✓	✓	✓	✓	✗	✓
(II) Client	✓*	✓	✓	✓	✓	✓	✓	✓
(III) Server	✓*	✓	✓	✗	✓	✓	✓	✓
(IV) User	✓*	✓	✓*	✗	✗	✗	✗	✓

*Observable in deteriorated data quality or abstracted form.
 Legend: Geo. = Geospatial, Iner. = Intertial, Phy. = Physiological.

4.3 VR Defenses

We highlight in Fig. 1 where the defenses can counter potential attacks and classify them based on five adapted categories. They consist of the two categories that De Guzman et al. [31] added to the primary three proposed by Roesner et al. [120], which are present in other privacy literature [48, 141, 155]. Given that many researchers highlighted the potential harm of deceptive immersive content [1, 18, 21, 36, 56, 91, 101, 144], we add a category for virtual content protection. Note that not all of these protections are related to *privacy* (§4.2), but also to *security* (i.e., measures to impede unauthorized data access [16]) and *safety* (i.e., measures to preserve the physical and mental well-being of users [36]). We highlight the following literature for guidance in security and safety attacks and protections: [31, 35, 39, 49, 69, 71, 102, 128, 134, 153]. We frame our SoK around attacks and defenses related to the *privacy* aspects of these defenses, mainly to input protection.

(I) Input Protection (Security & Privacy). Software that, e.g., perturbs [93] or abstracts [45] active (user) and passive (user’s surrounding environment) sensitive input information to prevent user privacy breaches. Additionally, systems should be secured against adversarial inputs that bypass detection (cyber attack in Fig. 1: 1A).

(II) Data Access Protection (Security & Privacy). Active and passive user inputs are stored, relayed and accessed to deliver user-consumable output. The corresponding privacy and security measures extensively overlap with other systems, which existing literature covers comprehensively [48, 55, 108, 141, 155].

(III) Output Protection (Security & Safety). Detecting and censoring [103] malicious manipulation of outputs can prevent security breaches like “clickjacking” [120] or physical harm, e.g., inducing collisions with obscured real objects [21, 144], VR sickness [21], and trigger epilepsy [10].

(IV) User Interaction Protection (Privacy & Safety). Privacy protections can enhance confidentiality (i.e., data is only revealed to selected entities [48]) in physical or virtual spaces shared by multiple interacting users, e.g., a private virtual enclave that other users cannot enter [44]. We add to this category safety measures such as invisible avatar barriers to avoid psychological harm from virtual harassment [18] or bullying [102].

(V) Device Protection (Security & Safety). Device security measures can implicitly protect users and data in the above defensive aspects, e.g., authentication prevents impersonation [76], and defend against cyberattacks targeting devices [102] and networks [51], and VR tracking system jamming [116], which could lead to physical harm.

(VI) Content Protection (Privacy & Safety). Safety measures such as age verification and content moderation can protect users against

immersive falsehoods, malicious advertisement [69], and inappropriate, unsolicited, and harmful content that may lead to mental harm, disinformation, or manipulation of opinions and ideals [18, 91]. The privacy concern involves detecting virtual content and environments nudging users to disclose sensitive information subtly, e.g., puzzles revealing health data [92].

5 TAXONOMY OF VR DATA & APPLICATIONS

Thanks to the sensor-generated data and the applications processing this information, users can experience VR. However, applications are also the gateway for adversaries to harvest sensitive user data and use such information against them. The following classifies and discusses the data attributes and the applications subject to our threat model.

5.1 VR Attributes

Method. Two researchers independently examined the 75 selected publications, focusing on the data attributes that were highlighted, attacked, or defended, and that pertained to users using the VR input devices enumerated in §2.1. Subsequently, they compared their findings and engaged in informed discussions to resolve conflicts. Would-be attackers can collect these attributes at different steps (data sources) of the VR information flow of § 4.1. Fig. 2 presents the resulting taxonomy of VR-derived data. We base our categorization on observable attribute classes and indicate which attributes or observations an attacker can directly capture from a data source (*primary*), deterministically derive from primary attributes (*secondary*), and infer from primary and secondary attributes employing ML or other learning procedures (*inferred*). Furthermore, we used the 75 publications to draw the connections between attributes. Specifically, the arrows in the figure represent the relationships that papers often explore, where one attribute is qualitatively discussed or quantitatively determined based on another. Each arrow signifies a direct connection made in the literature. Note that there might be other connections outside VR and new ones might arise in future work, e.g., deriving religious orientation or personality traits from VR inertial telemetry.

Geospatial & Inertial Telemetry. The position, orientation, and acceleration of body tracking devices over time reveal anthropometric measurements. Such measurements can be *direct* (body skeletal information such as arm-length and height [90]), *combined* to obtain further biometrics (e.g., wingspan [92]), or *compared* to draw relationships (ratios may reveal a user’s body asymmetries [93]). An attacker may also record kinesiological movements, which can reveal unique gestures [45, 128], or biometric movements [105] such as gait [125]. Additionally, the devices’ coordinates can map the play area’s boundaries, revealing its surface [92]. Even without full-body tracking devices, Winkler et al. [156] showed that reinforcement learning techniques could infer a full-body pose with telemetry from only an HMD, its IMUs, and hand-held controllers. Furthermore, Chen et al. [126] derived speech from the bone- and air-borne vibrations registered by an HMD’s IMU telemetry data, and Nair et al. [95] predicted with more than 80% accuracy traits like country of origin or foot size just with head and hand motion data. Note that hardware adversaries have a privileged position to observe device telemetry. In contrast, server and user adversaries

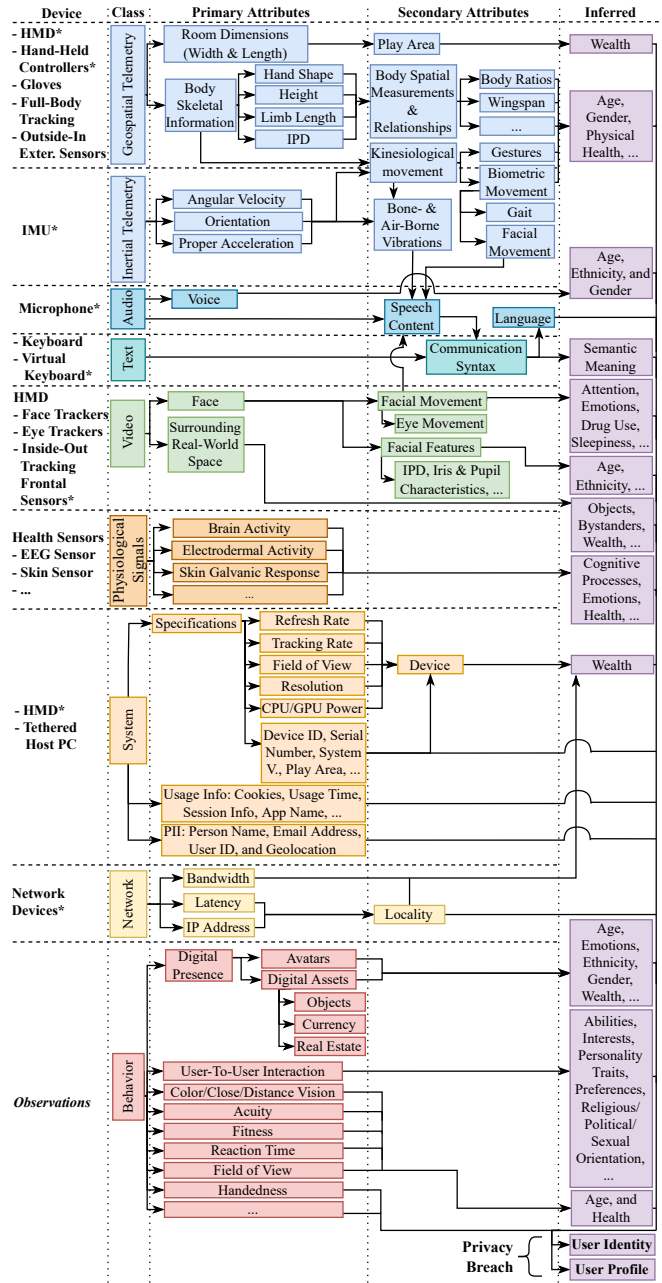


Figure 2: Taxonomy of VR data attributes. * Primary devices.

will experience degraded precision in their attribute estimations due to intermediate data processing, e.g., filtering and compression.

Audio & Text. Users can verbally interact with other users in virtual telepresence applications or give voice commands to their VR devices through a microphone [77]. Attackers can listen to vocalizations to fingerprint users based on vocal characteristics (e.g., frequency or accent) [92, 128] and profile them with communication semantics [40]. While voice biometrics may degrade along the data

flow, speech semantics are more robust and could remain vulnerable to user adversaries. Additionally, the messaging functionality enabled by physical or virtual keyboards operated with hand-held controllers or gloves increases the attack surface [7, 74, 124].

Video. HMD's face optical sensors can register and track eye and facial movements and features to render expressive photorealistic avatars [23]. However, the facial video feed can also serve to identify an individual (e.g., using IPD, or Iris, and pupil characteristics [22, 63]) or infer emotions [123, 162]. Notably, Krüger et al. [67] provided a comprehensive overview of the plethora of attributes that privileged adversaries can infer from eye tracking. Moreover, with expressive avatars, server and user adversaries could also learn other users' mental state. Additionally, while more prevalent in AR applications, the inside-out tracking frontal cameras of a VR HMD [83] also expose the real-world environment surrounding users, which can reveal sensitive information to hardware and client adversaries, such as personal objects [70, 160], the surrounding space type [28, 53], or bystanders [59, 60].

Physiological Signals. As health sensors like EEGs make their way into commercial-grade HMDs [99], the possibilities of VR (and privileged adversaries) expand dramatically. With these sensors, applications can adjust immersive experiences based on physiological signals that meet users' particular needs in real-time [5, 36, 106, 157] and can help users with rehabilitation treatments [4, 122, 128]. Such improvements, however, will also expose critically sensitive user information, such as physical and mental health conditions [36, 122, 157], behavior [5, 14, 137], language semantics [33, 138], and other sensitive PII like credit cards, PINs, and locations or persons known to the user [79].

System & Network. Adversaries can determine a user's VR device, host PC, network characteristics, and related internet session information [143]. Specifically, hardware and client adversaries can query system APIs to collect system specifications (e.g., tracking rate, resolution, etc.), and less privileged adversaries may devise attacks to gauge a target user's refresh rate without access to system APIs or user agents [92]. Notably, Trimananda et al. [143] captured the plethora of system information relayed to servers, which included all the above, in addition to PII like a person's name and usage information such as cookies or app names. On the one hand, while not specific to VR, as virtual telepresence applications rely on multiple servers to reduce perceived latency [151], attackers can observe network traffic to determine users' geolocation without an IP address. Altogether, these additional data points help adversaries fingerprint users to track them across internet VR sessions. Conversely, given that "System and Network" stand apart from applications and are well-understood domains, practitioners can employ conventional techniques to mitigate the corresponding vulnerabilities in the field of VR.

Behavior. Observing users' avatar likeness, expressed emotions, interactions and reactions to virtual stimuli from other avatars or virtual content can reveal various sensitive human characteristics [68, 123, 135]. In practice, malicious developers may carefully and inconspicuously deliver stimuli in a virtual experience to prompt the user to unconsciously reveal their reaction time, handedness, fitness level, visual and mental acuity, etc. [92]. Additionally, how a user chooses to represent their likeness as avatars, together with

the digital assets they own, can reveal information such as their demographics or wealth [36, 62]. Lastly, user-to-user interaction in social VR can lead to attackers directly spying on or engaging with the target user [44, 150]. The information required to meaningfully observe sensitive behavioral data is typically enough at each stage of the information flow [92].

Inferred Attributes. With the appropriate ML algorithm [43, 88, 110], the discussed attributes above can reveal demographics [3] and other related sensitive attributes such as emotions [123, 162], physical and mental health [67, 75], wealth, and political or sexual orientation or preferences over different users or products [44, 61], among others [20, 36]. Users may also unintentionally or voluntarily self-disclose such information or additional biographical data (e.g., age, home address, education, social status, work history, etc.) [131, 135], or be deceived by the application or other users to reveal inferable attributes [3]. Ultimately, adversaries can leverage the breadth of data to identify and profile users across VR applications.

There is an important dimension to consider when studying all these attribute types from a privacy perspective: *direct* and *indirect* leakage. Direct leakage is intentional, deriving from the application's primary objectives, like revealing user emotions in emotion-centric VR experiences. In contrast, indirect leakage, such as the unintended exposure of anthropometric data from inertial sensors, occurs inadvertently. The countermeasures for these two types differ greatly in implication and execution.

5.2 VR Applications

For decades, the gaming industry has advanced 3D graphics hardware and low-latency content delivery to create immersive, time-intensive online user experiences. Their expertise has pushed gaming to become the current dominant application in VR [11]. However, VR promises applications beyond entertainment: social life, education, healthcare, fitness, military training, architecture, retail, business, productivity (virtual offices), engineering, and manufacturing [24, 100]. Specifically, social VR has recently increased in popularity with titles such as *VRChat*, whereby users worldwide interact with each other in real-time [135].

Method. Among the 75 collected studies, only two classified VR applications based on the above target industries [40, 153]. In contrast, we provide an orthogonal categorization from a privacy standpoint inspired by [31, 92, 93] and our adversary, protection models, and taxonomy of attributes. Accordingly, we contemplate privacy risks in VR from three perspectives: (i) *adversarial*, (ii) *user protections*, and (iii) *data*. VR application developers may consider answering three questions:

(i) How much adversarial exposure could the application suffer? Fig. 3 shows the prevalence of hardware and client adversaries across all applications and the rise in privacy risks as users require servers to interact with others. While massively multi-user VR applications such as social VR are the most privacy-hostile environments, single-user applications are at least vulnerable to the VR firmware itself, as it may have direct network access to exfiltrate collected data from an application (e.g., Oculus Quest 2).

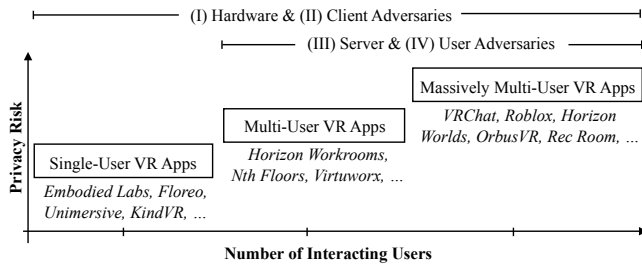


Figure 3: Privacy risk of VR applications as adversary exposure increases.

(ii) **How much privacy is the user willing to forgo using the application?** Some users are willing to expose all the information necessary to experience VR at its full immersive potential, while others are more reserved [35]. Hence, if protecting or opting out of specific data inputs is enabled (akin to internet cookies) [93], the privacy risks an application entails may vary from user to user. We suggest platform architects and application developers offer these protections and design their applications and games such that user experience for the privacy-conscious is not significantly deteriorated.

(iii) **How sensitive is the data handled by the application?** Most VR applications ingest *geospatial* and *inertial telemetry* and *audio*, and require a *system* and a *network* to join interactive experiences, where adversaries can extract *behavioral information*. These attribute classes form a privacy risk baseline. The application context raises the risks above this baseline, e.g., virtual health clinics, classrooms, and offices handle more PII and critically sensitive data than a game, e.g., *physiological signals*, *text* in homework or emails, and context-specific behavioral information such as attention to the lecturer or emotions during a meeting. The type of application and the context of its use add another critical dimension to the privacy risk depicted in Fig. 3.

6 VR ATTACKS

Method. Among the 75 collected studies, we found 35 attacks introducing explicit, offensive mechanisms (23) or methods that an attacker could leverage for adversarial purposes (11). For example, an attacker can leverage motion device authentication software to perform identification attacks across VR sessions. Two researchers iteratively discussed and systematically categorized these attacks in Table 2 (labeled with IDs A1 to A30). First, we employed a coding scheme derived from the threat model in §4.2 (i.e. attacker types), the attribute categorization from §5.1 (i.e. attribute types), integrated an evaluation metric, and the focus (VR or MR) of the publication. Following this schema, the two researchers individually extracted pertinent elements from each paper. Lastly, the two researchers compared their results and had an informed discussion to reconcile any discrepancies.

Where information was lacking, e.g., not all attacks had an explicit adversary model, we used our best judgment supported by the publications artefacts, e.g., the client was the most common adversary and studies such as A9, A11-12 developed an application. While the papers used several metrics, we included the most distinct

or prevalent metrics to benchmark the attacks. The two researchers designed the following six RQs to focus our findings: RQ1-2 discuss opportunities for attacks, and RQ3-6 give an overview of critical attacks and explore their viability and risk.

(RQ1) How can VR devices enable attack opportunities? According to the literature, the most accurate identification attacks rely on *HMDs* and *hand-held controllers* (position and acceleration) to capture kinesiological movements (A6-12), while eye trackers mainly have a supportive role (A10-12). Profiling attacks that predict critically sensitive information, namely emotions (e.g., arousal and stress levels), rely on *health sensors*. These attacks use devices such as EEGs (A26), electromyograms (EMG) (A23), and ECGs (A20, A26), but also blood pressure (A25), galvanic (A20, A25-27), thermal (A24-25), respiratory (A26-27), and photoplethysmographic (A9, A22) sensors. Particularly, accelerometer and EMG data are an effective combination for identifying users' reactions to virtual stimuli (A23), and EEGs are especially suitable for emotion prediction (A20). However, an attacker can also derive emotions from facial expressions reconstructed from a target ML model (A32). Lastly, we note that some VR devices and applications have security vulnerabilities [100], and current VR devices enable more accurate identification attacks than AR glasses (A11).

(RQ2) How can VR data attributes expose attack opportunities? The *geospatial telemetry* of HMDs and the hand-held controllers are a low-hanging fruit for adversaries. Attackers can simply measure biometrics like height and wingspan to uniquely identify a small set of users (A1, 30 users), and register unique motions such as pointing (which exhibit more identifiability than grabbing motions) and rich gestures from the dominant hand (A11). If the application is not sand-boxed, malicious code could exploit resource monitoring and allocation APIs from the game engine to derive voice commands or hand gestures (A33). Combined with *inertial telemetry*, an attacker can infer a user's full-body pose (even with avatars of different scales, A4), inferred typed words (A34), perform highly accurate identification attacks (A7, A8, A12), and infer age (A10). We found evidence that hand-eye coordination (A11) could provide stronger signals for identifiability than individual features, and eye-related patterns have considerable influence in gender prediction (A10), requiring the fusion with *eye movement video feeds*. In addition to proactive attacks, there is always the danger of unintentional or intentional self-disclosure (A30). With respect to the latter, user's *digital presence*, such as avatar likeness and assets, can disclose sensitive demographics, offering attack opportunities.

(RQ3) How invasive can VR attacks be? The malicious accumulation of user data through profiling and tracking across internet sessions can lead to surveillance advertisement [25, 69], price discrimination [47], cyber abuse [18], personal autonomy curtailment [36], and pushing political agendas [104], among others [1, 56]. These threats accentuate when adversaries can infer users' deep emotions, and reactions to stimuli [36], which, given how immersive experiences can be (A20), are more easily observable in VR. Accordingly, we find critically invasive the adversarial capability to design VR experiences that can adjust users' arousal (attention) (A26) and stress (A27) to the desired level, predict anxiety (A20), and recognize emotions (A32) and their causes (A16), in the name of personalized VR experiences. Other attacks include generative

Table 2: Systematization of VR attacks from collected papers.

ID	Name	Focus	Devices	Observable Attribute Classes										Metric						
				Geo. Tele.	Inertial Tele.	Text	Audio	Video	Physio. Signals	System	Network	Behavior	Profiling		Identification	Hardware	Client	Server	User	
A1	<i>MetaData</i> [92]	VR		•	-	-	•	-	-	•	•	•	•	•	•	-	•	•	-	Accuracy
A2	Malicious Design [3]	VR	N/A	-	-	-	-	-	-	-	-	-	-	-	-	•	•	-	-	N/A
A3	<i>MitR</i> [150]	VR	N/A	-	-	-	•	-	-	-	-	-	-	-	•	•	•	-	-	N/A
A4	<i>QuestSim</i> [†] [156]	VR		•	•	-	-	-	-	-	-	-	-	-	-	•	-	•	-	Geo. Errors
A5	<i>Face-Mic</i> [126]	MR		-	•	-	-	-	-	-	-	-	-	-	•	•	-	•	-	Accuracy
A6	<i>GaitLock</i> [†] [125]	MR		-	•	-	-	-	-	-	-	-	-	-	-	•	•	-	-	Accuracy
A7	Movement Biometrics [†] [91]	VR		•	•	-	-	-	-	-	-	-	-	-	-	•	•	-	-	EER
A8	Movement Biometrics [90]	VR		•	•	-	-	-	-	-	-	-	-	-	-	•	-	•	-	Accuracy
A9	Movement Biometrics [†] [73]	VR		•	•	-	-	-	-	-	-	-	-	-	-	•	-	•	-	Accuracy
A10	Movement Biometrics [142]	MR		•	•	-	-	•	-	-	-	-	-	-	•	•	-	•	-	F1-Score
A11	Movement Biometrics [†] [111]	VR		•	•	-	-	•	-	-	-	-	-	-	•	-	•	-	-	Accuracy
A12	<i>BioMove</i> [†] [105]	VR		•	•	-	-	•	-	-	-	-	-	-	•	-	•	-	-	Accuracy
A13	Eye Tracking [‡] [28]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	Accuracy
A14	Iris Identification [‡] [22]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	Accuracy
A15	<i>Kaleido</i> [‡] [72]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	-	•	-	F1-Score
A16	<i>EMOShip</i> [†] [162]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	F1-Score
A17	Spatial Recognition [53]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	MER
A18	Spatial Recognition [‡] [70]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	F1-Score
A19	<i>SafeMR</i> [‡] [30]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	•	-	Accuracy
A20	<i>Vreed</i> [†] [137]	MR		-	-	-	-	•	•	-	-	-	-	-	•	-	•	-	-	Signal Statistics
A21	<i>Galea</i> [†] [14]	VR		-	-	-	-	•	•	-	-	-	-	-	•	-	•	-	-	Signal Statistics
A22	Signal Processing [†] [80]	VR		-	•	-	-	-	•	-	-	-	-	-	•	-	-	•	-	EER
A23	Signal Processing [†] [5]	VR		-	•	-	-	-	•	-	-	-	-	-	•	-	-	•	-	Signal Peaks
A24	Signal Processing [†] [106]	VR		-	-	-	-	-	•	-	-	-	-	-	•	-	•	-	-	Signal Stability
A25	Signal Processing [†] [122]	VR		-	-	-	-	-	•	-	-	-	-	-	•	-	-	•	-	Signal Statistics
A26	Signal Processing [†] [157]	VR		-	-	-	-	-	•	-	-	-	-	-	•	-	-	•	-	Accuracy
A27	Signal Processing [†] [58]	VR		-	-	-	-	-	•	-	-	-	-	-	•	-	-	•	-	Signal Statistics
A28	Virtual Typing [74]	VR		•	•	•	-	•	-	-	-	-	-	-	•	•	-	-	•	Accuracy
A29	<i>VR-Spy</i> [7]	VR		-	-	•	-	-	-	-	-	-	-	-	•	•	-	-	•	Accuracy
A30	Self-Disclosure [†] [135]	MR	N/A	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	•	N/A
A31	Movement Biometrics [94]	VR		•	-	-	-	-	•	-	-	-	-	-	•	•	-	-	•	Accuracy
A32	Facial Reconstruction [123]	VR		-	-	-	-	•	-	-	-	-	-	-	•	-	-	-	•	Accuracy
A33	Side Channels [161]	MR		-	-	-	-	-	-	•	-	-	-	-	•	-	•	-	-	F1-Score
A34	Movement Biometrics [130]	MR		•	•	-	-	-	-	-	-	-	-	-	•	-	-	•	-	Top-k Accuracy
A35	Movement Biometrics [69]	VR		•	•	-	-	-	-	-	-	-	-	-	•	-	-	-	•	Accuracy

[†]An attacker can leverage the defense/mechanism for adversarial purposes. [‡]Although the study is defense focused, there is an adversarial component.
Names: Names in italics correspond to the authors' selected title, otherwise, it is a descriptive name.
VR Device: = HMD, = Eye Trackers, = Inside-Out Tracking Optical Sensors, = Outside-In Tracking Optical Sensors, = IMU Orientation, = IMU Velocity, = Hand-Held Controllers, = Microphone, = Tethered PC, = Network Devices, = Health Sensor, N/A = Not Applicable.
Abbreviations: Geo. = Geospatial, Tele. = Telemetry, Physio. = Physiological, PB = Privacy breach, MER = Mean error rate, EER = Equal error rates.

AI producing deepfakes to create false memories (A2), virtual content nudging users to unintentionally leak information (A1), while others pervasively deduced text, via automated external observers that infer user-typed passwords from arm movements (A28), or by using channel state information of WiFi signals to infer unique gesture patterns and the corresponding keystrokes (A29).

(RQ4) How practical are privacy attacks? *User* adversarial attacks are easy to execute (A3, A30), as they can, at a minimum, join a VR session and social engineer information from users. These attacks aggravate when exploitable bugs allow, e.g., invisible avatars

(A3). In contrast, attacks relying on physiological signals still require researchers to enhance their own VR *hardware* to register ECGs (A22), electrodermal and muscular activity (A21), or photoplethysmographic data (A24), indicating a lack of maturity of such attacks. Additionally, given current practices and the low establishment of VR privacy standards and enforcement [143], an adversary in control of the *client* application or the *server* running multi-user functionality could put into practice the associated attacks the researchers have demonstrated. Most critically, they can infer in a few minutes more than 25 attributes, including sensitive demographics

(A1), hide malicious operations that collect or infer information from an otherwise honest application (A18), run emotion detection using video feeds (A16), and run authentication algorithms to tag users (A6-A12). Thus, it is paramount that VR platforms add permission systems to limit the scope of these potential adversaries. This real-world mitigation emphasizes the shared responsibility between developers and platform designers in safeguarding user privacy.

(RQ5) How effective are privacy attacks? Based on the literature, *identification* attacks targeting kinesiological movements are highly accurate. Particularly, the most effective *identification* attack targets gait and relies on dynamic time warping and sparse representation classifiers to achieve an accuracy of 98% using only IMUs (A6, experiments with 20 users). Others reach such accuracy by additionally feeding translational movements to a CNN (A7, 41 users), in addition to eye tracking fed into kNN and SVM classifiers (A12, 15 users). In contrast, the most robust attack, i.e., tested with 50,000+ participants, achieved an accuracy of 94.33% feeding a gradient boosting decision tree [66] the positions of the HMDs and controllers of a 100-s interval (A31). Techniques that improve the resulting accuracy comprise normalization of height and arm lengths (A9), and smoothing methods for pre-processing (A7).

Regarding *profiling*, one particularly effective and broad attack (A1, 30 users) accurately measured from an HMD and its hand-held controllers multiple primary attributes such as height, wingspan, handedness, and interpupillary distance, among others. With this data, A1 inferred gender, age, and ethnicity with close to 100% accuracy, using variations of SVM and random forest classifiers. Furthermore, the most effective emotion profiling attacks in the literature achieved an accuracy of 80-90%. They used ECG's signal amplitude and eye tracking data (A20, 12 users) or primarily EEG (A26, 12 users) as inputs to an SVM, or fed facial and surroundings video to tailored ML pipelines (A16, 20 users).

(RQ6) How these VR attacks conform with a well-established threat model? To conclude, we briefly examine these attacks from the perspective of a non-VR specific, highly-cited threat model: Lindunn [34], which guides the systematic elicitation and mitigation of privacy threats in software architectures. All the successful attacks in Table 2 imply users' *Content Unawareness* and system's *Policy and Consent Non-Compliance* because the users are unaware of the hidden malicious operations that exploit the permissions granted and advertised by the VR application. Moreover, these attacks would not be possible without incurring a *Linkability* or *Detectability* threat, as a successful attack must correctly link and assess the existence of a user and an attribute. Furthermore, attacks profiling and identifying users map to the *Information Disclosure* and *Identifiability* threats, respectively. Furthermore, although it was not explicitly specified, some of these attacks could also *Non-Repudiate*, i.e., adversaries could show proof of the user's private virtual activities or attributes (A1, A3, A17).

7 VR DEFENSES

Method. Using an information extraction method akin to our approach with VR attacks, we categorized the 35 recognized defenses (labeled as D1 to D35) based on the defense model detailed in §4.3 and the attribute classification from §5.1. Table 3 systematically categorizes the defenses based on the corresponding papers. Similarly,

we designed four RQs. While RQ1-6 draw researcher attention to areas where attacks usually excel, RQ7-10 highlight where defenses do not necessarily do: we classify defenses (RQ7) to provide a frame, and highlight usability (RQ8), limitations (RQ9), and practicality (RQ10) as focus areas for researchers—§8 discusses opportunities.

(RQ7) What are the types of defensive mechanisms?

(i) *Perturbation*. Some works provide provable privacy guarantees by adding noise (differential privacy (DP) [41]) to geospatial or eye tracking data (e.g., D1, D4), while others blur (e.g., D11, D18) or mask (e.g., D3, D17) regions of a video like facial features, sensitive objects, and bystanders.

(ii) *Information abstraction*. Software that extracts key features from the surrounding space (e.g., surfaces, D16) or shares only the events triggered by sensitive inputs (e.g., unique gestures, D25).

(iii) *Recognizers*. Automated deepfake detection (D32), and middleware that detects and warns the user of sensitive surrounding objects and bystanders (D15).

(iv) *Static & Dynamic Analyzers*. Detection of application vulnerabilities that could lead to, e.g., unauthorized access to a private VR room (D33), and malware that, e.g., detects and exfiltrates sensitive surrounding objects (D13).

(v) *Platform features*. These solutions comprise mainly user-interaction protections. The primary examples include virtual (and physical) private enclaves that only authorized users can trespass (D31). Moreover, other defenses focus on confusing adversaries, e.g., with avatar clones dispersed across multiple VR applications, teleportation to new virtual locations, private copies of the virtual public environment, and platform-generated non-identifiable or invisible avatars (D30). Furthermore, platforms could include embedded voice modulators and social media privacy settings, whereby, e.g., only friends could see one's avatar (D32).

(vi) *Authentication*. Biometric movement recognition for logging into a VR device (D26-29).

(RQ8) How defenses balance usability and privacy? Usability is critical for immersion in VR applications; thus, researchers design utility metrics to assess the loss of usability when the user enables privacy protections. Aspects that impact usability are battery *energy consumption* (D14, D17), *latency* (D2, D4), and *playability* (D1, D4, D9, D14), i.e., how enjoyable or productive a VR experience is. Approaches that help to minimize device *energy consumption* are a tethered PC, offloading computation to the cloud (although bandwidth may become a challenge, D17), and sharing processing resources like object detection with other applications, which also reduces *latency* (D14). VR protections can decrease *playability* if the defense perturbs data, which is measurable primarily with metrics such as game scores (D1), subjective enjoyment (D4), attentiveness, comfort (D9), naturalness (D23), and accuracy loss (D1, D4, D16), among others. Additionally, enabling users to manage their privacy makes protections more usable (D31). Hence, applications empower users by giving them a choice (and the responsibility) to switch protections on depending on the context (D11), select their privacy strength with modulators like sliders (D1, D4), and providing visual prompts that communicate the impact of applications accessing user data (D15). In the background, these choices change the parameters quantifying privacy (e.g., ϵ in DP), values which we suggest setting empirically ex-ante (D1-2, D4, D5).

Table 3: Systematization of VR defenses from collected papers.

ID	Name	Focus	Devices	Defended Attribute Classes											BP	Protection	Metric							
				Geo. Tele.	Inertial Tele.	Text	Audio	Video	Physio. Signals	System	Network	Behavior	Profiling	Identification				Input	Data Access	Output*	User Interaction	Device	Content	
D1	<i>MetaGuard</i> [93]	VR		•	-	-	•	-	-	-	•	-	-	-	•	•	•	-	-	-	-	-	-	Accuracy
D2	Eye Tracking [28]	MR		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	F1-Score
D3	Iris De-Identification [22]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Accuracy
D4	Kaleido [72]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	Accuracy
D5	Eye Tracking [131]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	Accuracy
D6	Eye Tracking [154]	VR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	QoE
D7	Eye Tracking [27]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	Accuracy
D8	Eye Tracking [75]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	-	•	-	-	-	-	-	CC, MSE
D9	Eye Tracking [62]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	CRR
D10	Eye Tracking [17]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	CC, NMSE
D11	<i>EyeVEIL</i> [63]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Accuracy
D12	<i>PrivacEye</i> [133]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	-	-	-	-	Accuracy
D13	Spatial Recognition [70]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	-	•	-	-	-	-	-	F1-Score
D14	<i>SafeMR</i> [30]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	-	-	-	-	Accuracy
D15	OS Support [26, 59]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	-	-	-	-	-	-	-	FN, FP
D16	Spatial Recognition [29, 32]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	-	-	-	-	Accuracy
D17	<i>PlaceAvoider</i> [139]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	-	-	-	-	Accuracy
D18	<i>Darkly</i> [60]	MR		-	-	-	-	•	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	# Breaches
D19	Spatial Recognition [160]	MR		-	-	-	-	•	-	-	-	-	-	-	•	-	•	-	-	-	-	-	-	Accuracy
D20	Spatial Recognition [136]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	N/A
D21	<i>OpenFace</i> [152]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Accuracy
D22	<i>GARP-Face</i> [38]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Accuracy
D23	GAN-Based Defense [19]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Naturalness
D24	GAN-Based Defense [159]	MR		-	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Accuracy
D25	<i>Prepose</i> [45]	MR		•	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	Expression
D26	Movement Biometrics [73]	VR		•	•	-	-	-	-	-	-	-	-	-	-	†	-	-	-	-	•	-	-	Accuracy
D27	Movement Biometrics [111]	VR		•	•	-	-	•	-	-	-	-	-	-	-	†	-	-	-	-	•	-	-	Accuracy
D28	<i>GaitLock</i> [125]	MR		-	-	•	-	-	-	-	-	-	-	-	-	†	-	-	-	-	•	-	-	Accuracy
D29	<i>BioMove</i> [105]	VR		•	•	-	-	•	-	-	-	-	-	-	-	†	-	-	-	-	•	-	-	Accuracy
D30	Digital Presence [44]	MR	N/A	-	-	-	-	-	-	-	-	-	-	-	•	•	-	-	-	-	•	-	•	N/A
D31	<i>SecSpace</i> [117]	MR	N/A	-	-	-	-	-	-	-	-	-	-	-	-	•	•	-	-	-	-	-	-	N/A
D32	Design Defense [3]‡	VR	N/A	-	-	-	-	-	-	-	-	-	-	-	-	•	-	-	-	-	•	-	•	N/A
D33	<i>MitR Defense</i> [150]‡	VR	N/A	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	•	-	-	FN, FP
D34	Self-Disclosure Defense [78]	VR	N/A	-	-	-	•	-	-	-	-	-	-	-	•	•	-	-	-	-	•	-	-	N/A
D35	<i>ReconViguration</i> [124]	VR		-	-	•	-	-	-	-	-	-	-	-	•	•	•	-	-	-	-	-	-	Error Rate

‡Although the study is attack focused, there is a defensive component. † Authentication protection (as opposed to identification protection).

*Output safety and security attacks and defenses are covered in dedicated literature [10, 21, 103, 120, 144].

Names: Names in italics correspond to the authors' selected title, otherwise, it is a descriptive name.

VR Device: = HMD, = Eye Trackers, = Inside-Out Tracking Optical Sensors, = Outside-In Tracking Optical Sensors, = IMU Orientation, = IMU Velocity, = Hand-Held Controllers, = Microphone, = Network Devices, = Physical keyboard; N/A = Not Applicable/Available.

Abbreviations: Geo. = Geospatial, Tele. = Telemetry, Physio. = Physiological, BP = Breach prevention, QoE = Quality of experience, CC = Correlation coefficient, (N)MSE = (Normalized) Mean square error, CRR = Correct recognition rate, FN = False negative, FP = False positive.

(RQ9) What are the limitations in privacy defenses? Based on the literature, we suggest several key improvement areas: (i) Update *perturbation* protections that use generative adversarial networks, considering the rapid advancements in generative AI [37], as the latest study available is almost four years old (D24). (ii) Enhance biometric-movement authentication to cover activities beyond walking, like running (D28). (iii) Expand provable privacy guarantees to eye tracking (D4, D8, D10) and geospatial (D1) time series. (iv) Increase research on user privacy preferences in VR to contextualize protections, (v) improve permission structures of VR operating systems (D13), and (vi) develop a standard vetting

program to verify VR library functions (D16). However, the closed-source nature of VR systems [84] may limit defensive deployments. **(RQ10) How practical and effective are privacy defenses?** Researchers typically implement defenses as middleware (D15-16, D18) that pre-processes data before a potentially malicious application ingests it, or as an easy-to-install plugin within the application (D1). In terms of effectiveness, the latter would only defend against *server* and *user* adversaries. Further, unless the implementation is at the firmware level, users have no protection against a *hardware* adversary. Additionally, we identified in the literature promising prototypical defenses that significantly reduced the accuracy of

identification attacks based primarily on geospatial (D1) and eye tracking (D4) data to random-guess levels. Moreover, we found works demonstrating substantial accuracy degradation in *profiling* attacks at reasonable privacy levels (DP with $3 \leq \epsilon \leq 6$). Notable examples include *gender* inference based on eye tracking, whose accuracy D5 reduced to random guessing, and deriving *age* primarily from geospatial telemetry, which D1 reduced by 58.25%.

8 VR PRIVACY OPPORTUNITIES

This section complements the taxonomies by exploring privacy practices and opportunities based on a quantitative analysis of attacks and defenses (Tables 6 and 7 of Appendix E), and a qualitative examination of research gaps (Table 4), and in practice.

8.1 Quantitative Analysis

Our primary objective is not to verify the results of the studies, instead, it illustrates a possible method of ranking attributes based on their protectability, with an intent to inspire researchers to routinely incorporate this analysis in their investigations. To accomplish this, we searched for pairs of complementary open-source works (one attack and one defense) that considered a wide range of granular attributes that could be ranked in terms of privacy risk and defensibility. Within the limited pool of only four open-source defense-focused works, only *MetaGuard* (D1) [97] satisfied our criteria, primarily due to its unambiguous coupling with an attack study ((A1) *MetaData* [96]).

Evaluation Method. Tables 6 and 7 present the results of this analysis, which consists of three steps:

(i) *Risk.* We ran a PCA with Azure ML [87] over the anonymized ground truth of the participants of the *MetaData* study to calculate the amount of variability explained in PC1 by each attribute (e.g., height, or wingspan) for each inferred data point (e.g., gender, or age). Summing attribute contributions yields a summary statistic of the risk of attribute leakage, representing the information adversaries could obtain from their observations (Table 6).

(ii) *Weighted Mean Degradation.* For this metric, we relied on the anonymized frame-by-frame telemetry data of the 30 participants of the *MetaData* study, which we used to replicate the attacks with nearly identical accuracy. The replicated attacks measured the sensitive target attributes. Consecutively, we repeated the attacks with *MetaGuard* enabled at three privacy levels: low, medium, and high, to measure the degraded new attack accuracy. With these results, we performed a weighted average on the degraded attack accuracy at different privacy protection levels to reveal the attributes easiest to protect, i.e., with the highest accuracy degradation (Table 7).

(iii) *Opportunity.* Finally, we ordered the attributes with the highest accuracy degradation and risk in Table 7, highlighting the most sensitive and easiest-to-protect attributes.

Opportunities. We suggest privacy practitioners to prioritize deploying differential privacy defenses that protect *room size*, *height*, and *interpupillary distance* (IPD) in their devices and applications, as they show the highest leakage risk and sensitivity to noise.

Ethical Considerations This SoK does not contain original data collected from human subjects. We replicated prior studies employing anonymous data collected directly from the authors of those studies or publicly available online repositories. We verified that

OHRP-registered institutional ethics review boards processed and approved those prior studies and considered them non-deceptive. Additionally, those studies' informed consent documents included permission to re-use collected data for follow-up research work, and we handled such data rigorously according to the corresponding original consent documentation requirements.

8.2 Min-Set Coverage of Defenses & Attacks

Method. We selected the most comprehensive attacks per attribute class and mapped them to the most fitting defenses in Table 4, considering identification and profiling privacy breaches and highlighting privacy opportunities.

On the one hand, While many attacks possess corresponding defenses, several defenses don't comprehensively address their associated attack vectors. For instance, defenses for audio identification have seen more defenses outside of VR [98], (D35) ReconViguration overlooks its counter (A29) *VR-Spy*, which utilizes WiFi signal's channel state information to deduce distinct gesture patterns. Additionally, (D1) *MetaGuard* falls short against malicious content used by (A1) *MetaData* to subtly gauge users' cognitive functions. Given these gaps, we urge the research community to revisit and fortify these defense strategies.

On the other hand, Table 4 underscores significant research gaps in the VR threats and protections landscape, presenting new avenues for scientific exploration. Although there is orthogonal research on topics such as brain waves and physiological signals [33, 79, 138], these areas remain largely unexamined from an adversarial or defensive standpoint within the specific context of VR. Additionally, defenses against utilizing geospatial and inertial data for identification and profiling (e.g., A5-6, A10, A12) are conspicuously lacking. Further investigation into the impact of defenses on usability, such as the memory-related effects of blurring bystanders, could also provide valuable insights [42].

8.3 VR Defenses in Practice

Method. We examined the 58 studies on attacks and defenses to find open-source implementations of their proposed mechanisms, and briefly explored privacy threats and protections in the industry.

Of the repositories analyzed, only 21% were functionally operative: 4 for defenses (D1, D16, D25, D33) and 8 for attacks (A1, A7-8, A16-17, A21, A27, A31). For instance, while *MetaGuard* (D1) introduced the first VR "incognito mode" as a Unity plugin available for any VR app via MelonLoader [81], only *Prepose* (D25) was officially affiliated with Microsoft—though without clear evidence of production use. Notably, Bigscreen incorporated suggestions from *MitR Defense* (D33) to address their privacy issues. Conversely, *EMOShip* (A16) is integrated into Pupil Labs' tech stack. Overall, we observe little transfer of privacy research into the VR industry. Additionally, the conclusions from the 2022 evaluation study *OVRseen* [143] indicated a significant lack of privacy measures in commercial-grade applications: 70% of VR data flows from the most widely adopted VR device were not appropriately disclosed, and 69% of them were used for purposes different to the core apps' functionality. Hence, we encourage researchers to systematically open-source their code and engage with companies to bring privacy protections to consumers.

On the industry side, several indicators are forming a trend not conducive to enhancing privacy. Companies have not shipped a “VR incognito mode” to avoid user tracking across VR applications, some developers ignore their own privacy policies [143], consumers need to pay extra to sign in to their VR headsets without a social media account [114], the patents of a major VR company reveal how face tracking will help with personalized advertising in future metaverse applications [140], policy updates trend towards more data collection [143], VR devices and applications are shown to be vulnerable to exploits [100], and companies can ban plugins [57] that could help with privacy, safety, and user disabilities.

Nonetheless, one of the most advanced, commercial-grade VR device released on October 2022 [84] has proposed a set of privacy features that update the industry standards [85]: (i) monitoring features are turned off by default, (ii) tracking is paused at headset removal, (iii) cameras and microphone are turned off during headset’s sleep mode, (iv) raw images are processed, stored, and deleted locally, (v) the extracted features are not used to identify users, and (vi) external lights on the headset signal bystanders that outwards cameras might record them. However, there exist caveats, e.g., the company’s eye-tracking notice [82] indicate that abstracted facial tracking information could be stored and processed by servers (e.g., potentially for psychographic profiling), abstracted gaze data can be shared with third parties (where the data is subject to their own privacy policies), and dark patterns are prevalent, namely, “Enable” buttons are more highlighted than the “Not Now” option. More advanced VR headsets, released in June 2023 [6], also emphasized the importance of user privacy by processing both optic biometrics for identification and eye tracking data locally. This data is stored in encrypted form, and applications use system APIs for producing spatial experiences without accessing the raw sensor information.

Overall, given these industry and research privacy gaps, there are numerous opportunities for academics and practitioners to improve the state-of-the-art in privacy-enhancing VR systems, e.g., adding features such as end-to-end encryption for biometrics like eye-tracking, control programs to verify third party apps, more granular user privacy options, and abstracting biometric time-series to events [12].

8.4 VR Incognito Mode

To develop a comprehensive “VR incognito mode,” or rather, “ecosystem” that complies with the defense model of § 4.3, we suggest practitioners combine the privacy-forward features outlined earlier and those listed in Table 4. Ideally embedded at the firmware level by VR device manufacturers (where applicable), this mode should impartially apply its privacy protections, without favoring its developer or device manufacturer. This approach would establish a robust privacy framework for VR:

- (i) *Privacy-Aware Hardware Design*: Devices should seamlessly incorporate privacy features, such as auto-pausing tracking and sensor recordings—especially cameras and microphones—when not in use and notifying bystanders about possible recordings.
- (ii) *Privacy-Forward OS Design*. The VR OS should be inherently privacy-centric. For example, it ought to vet third-party apps for foundational privacy and security standards, implement app sandboxing, convert raw sensor data into event-driven inputs, and make

them accessible via permission-controlled APIs, thereby preventing continuous access to raw data. The OS ought to prioritize in-device data processing and encryption.

(iii) *On-Device Processing*: While this aspect is embedded within the OS design, its criticality mandates separate attention. Prioritize local data processing and encryption to ensure sensitive user interactions are not transmitted or stored externally, including identification.

(iv) *Browsing & Communication Safeguards*: Prevent the retention of user interactions by omitting browsing histories, cookies, and other digital fingerprints in VR. Introduce end-to-end encryption for data required to interact with other users.

(v) *Privacy-Enhancing Technologies*: Incorporating defenses from Table 4 against identification and profiling attacks, while maintaining the quality of the VR experience, would empower users.

(vi) *Intuitive Privacy Settings*: Design an easily navigable interface for users to effortlessly and granularly adjust their privacy settings. Avoid dark patterns and prioritize flexible, user-focused privacy management. By default, pervasive features like monitoring should be turned off.

(vii) *User Data Control*: Such settings should provide users with transparency and authority over shared data, ensuring they maintain ownership of their digital identity.

Table 4: Min-Set coverage of attacks and defenses.

Class	Privacy Attack	Privacy Defense
Identification		
Geospatial Telemetry	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i>
	(A12) <i>BioMove</i> [†]	★
Inertial Telemetry	(A6) <i>GaitLock</i> [†]	★
	(A5) <i>Face-Mic</i>	★
Text	(A28) <i>Virtual Typing</i>	(D35) <i>ReconViguation</i>
	(A29) <i>VR-Spy</i>	(D35) <i>ReconViguation</i> [*] , VPN, Tor, Proxies, End-To-End Encryption.
Audio	(A3) <i>MitR</i>	(D33) <i>MitR Defense</i> [‡]
	Speech Recognition	Voice Modulation [78, 98]
Video	(A13) <i>Eye Tracking</i> [‡]	(D2) <i>Kaleido</i> [*]
	(A14) <i>Iris Identifi.</i> [‡]	(D11) <i>EyeVEIL</i>
Physio. Signals	★	★
System	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i> [*]
Network	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i> [*] , VPN, Tor, Proxies.
Profiling		
Geospatial Telemetry	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i> [*]
	(A10) <i>Movement Bio.</i>	★
Inertial Telemetry	(A5) <i>Face-Mic</i>	★
	(A10) <i>Movement Bio.</i>	★
Text	<i>Ditto</i>	
Audio	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i> [*]
	(A3) <i>MitR</i>	(D33) <i>MitR Defense</i> ^{‡*}
Video	(A16) <i>EMOShip</i> [†]	(D12) <i>Kaleido</i> [*]
	(A17) <i>Spatial Recog.</i>	(D16) <i>Spatial Re.</i>
	(A18) <i>Spatial Recog.</i> [‡]	(D13) <i>Spatial Re.</i>
Physio. Signals	(A20) <i>Vreed</i> [†]	★
	(A22) <i>Signal Proces.</i> [†]	★
Behavior	(A1) <i>MetaData</i>	(D1) <i>MetaGuard</i> [*]
	(A3) <i>MitR</i>	(D33) <i>MitR Defense</i> ^{‡*}
	(A2) <i>Malicious Design</i>	(D32) <i>Design Defense</i> ^{‡*}

[†] An attacker can leverage the associated defense/mechanism for adversarial purposes. [‡] While the study is attack/defense focused, there is a defensive/adversarial component. ^{*} Privacy opportunity.

9 DISCUSSION & FUTURE WORK

We distill key findings (KF) and future work (FW) from studying the 75 selected publications and our results:

(KF1) *There is a fundamental imbalance in the research and deployment of offensive and defensive VR research.* The emphasis on defensive publications (see Table 5) could reflect the field’s emerging stage, underscored by the sparse open-source contributions. Additionally, given that attacks in web technologies have developed over decades, many may overlap with VR [98], suggesting that the real gap lies in crafting VR-specific privacy defenses. Over such decades, an extensive set of countermeasures has also developed and become widely adopted in the field of web privacy, suggesting a direction for the nascent VR privacy research.

While most reviewed papers focused on defensive techniques, none were deployed. This indicates limited knowledge transfer to industry, as evident by the absence of open-source code in academic papers. In contrast, we have so far seen a plethora of device vulnerabilities [100], an increased data collection and privacy policy disregard [143], along with one deployed academic offensive technique (A16).

(KF2) *Only a few defense studies (17%) provided provable privacy guarantees.* Provable privacy appeared in eye tracking studies (D2, D4-5, D8, D10) and spatial telemetry (D1). However, provable privacy is still uncommon in VR, possibly due to its relative immaturity.

(KF3) *VR authentication mechanisms should remain in the device.* Selected studies suggest that movement biometrics can be used for identification attacks. While we acknowledge the merits of privacy-preserving server-side authentication, on-device authentication—akin to mobile phone facial ID verifications—offers fewer potential attack vectors since the data does not leave the device. Especially, when continuous authentication is a requisite (e.g., during exams taken in VR).

(KF4) *There is a lack of hardware-level privacy defenses.* E.g., practitioners could use (D1) *MetaGuard* at the firmware level or execute a signed function in a trusted execution partition instead of relying on performance variations to detect application misbehavior (D13).

(KF5) *Attack benchmarks should use appropriate metrics, and defense proposals should generally include usability and performance studies.* We encourage researchers to add *F1-Scores* or *equal error rates*, as false positives and negatives are essential for security and privacy, and measure performance degradation in, e.g., execution time, battery consumption (D14, D25), and usability (D1, D9).

(KF6) *We identified the most dangerous attacks.* Based on the literature, (A31) [94] and (A35) [95] are the most effective and practical identification and (broadest) profiling attacks, respectively.

(KF7) *There is an asymmetry in the focus on attacks and defenses across various data attributes.* Text, audio, system, and network attributes are less explored, suggesting that solutions might exist in other research fields. This highlights the need to integrate insights from other disciplines into VR.

(KF8) *The majority of attack research targets VR, while defense research predominantly considers MR.* Based on the distinctions outlined in Appendix A, VR-specific attacks accounted for 59% and were particularly concentrated on physiological signals and motion telemetry. In contrast, 74% of defenses catered to MR, encompassing AR. This emphasis on MR in defense research is likely due to the

presence of frontal cameras and eye trackers in both VR and AR devices, with the majority of selected defense publications centering on these components.

Beyond the privacy opportunities highlighted in §8 and the limitations of RQ9, we encourage researchers to:

(FW1) *Design protections for spatial and inertial telemetry against adversarial interference.* The majority of surveyed defenses focused on protecting information extracted from video feeds. Moving forward, we hope to see defensive research fill the remaining gap.

(FW2) *Explore attacks of VR-native stimuli.* While some reviews discussed the safety dangers of maliciously manipulating video output [10, 21, 144], none investigated if such dangers apply to audio, stereoscopic vision, haptic feedback, or other VR-specific outputs. There is a lack of defenses against such risks.

(FW3) *Develop concrete countermeasures against malicious content design.* Mitigating the ability of adversaries to gain information or influence users by manipulating the immersive VR environment is amongst the most difficult open problems in VR. Achieving an appropriate balance between flexibility and consumer protection in VR environment design remains a significant outstanding challenge.

(FW4) *Resume research in privacy protections for VR user interaction.* The deeply immersive nature of VR makes social engineering a salient threat, but many extant studies on this subject are outdated by more than 10 years [31].

(FW5) *Study the inference of health conditions based on physiological signals in VR.* While studies have inferred emotions and arousal in VR, they underscore the potential for adversaries to discern neurological issues, physical disabilities, addictions, and health conditions like asthma. This emphasizes the urgency of defensive research, given the current emphasis on attacks and complete lack of defenses for physiological signals.

(FW6) *Explore the use of trusted execution environments (TEEs) in VR.* At server or client level, TEEs could enhance privacy beyond surveyed defenses. Deploying TEEs for GPUs might open new privacy avenues.

(FW7) *Systematize users’ perceptions of VR privacy.* Gaining insights into users’ perspectives on privacy in VR can lead to more effective privacy protection strategies. Additionally, these insights could aid in improving VR interface design, creating a safe environment where users can confidently interact, which could in turn bolster VR adoption and trust.

(FW8) *Distinguishing direct and indirect information leakage in VR.* Recognizing this distinction is paramount as the countermeasures and resulting user experience can vary significantly between the two. Developers need to design privacy measures accordingly, while ensuring users are informed about potential VR data disclosures.

10 CONCLUSION

In this SoK, we present a threat and defense framework for data privacy in VR and outline privacy opportunities for practitioners. Despite more defense proposals than attacks in literature, existing defenses are not exhaustive, some are missing, and most remain undeployed. The rise of data-hungry companies and pervasive data collection in VR highlights the need for increased cross-collaboration between industry and academia. Our frameworks and taxonomies aim to provide a foundation for future collaborations and research on metaverse privacy issues.

REFERENCES

- [1] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics Emerging: the Story of Privacy and Security Perceptions in Virtual Reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 427–442.
- [2] Rachel Albert, Anjul Patney, David Luebke, and JooHwan Kim. 2017. Latency Requirements for Foveated Rendering in Virtual Reality. *ACM Trans. Appl. Percept.* 14, 4, Article 25 (sep 2017), 13 pages. <https://doi.org/10.1145/3127589>
- [3] Nadisha-Marie Aliman and Leon Kester. 2020. Malicious Design in AIVR, Falsehood and Cybersecurity-oriented Immersive Defenses. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. 130–137. <https://doi.org/10.1109/AIVR50618.2020.00031>
- [4] Amelia Virtual Care. 2022. Virtual Reality Solution for Mental Health Professionals. <https://ameliavirtualcare.com/virtual-reality-solution-psychology/>. Online; accessed 22 Sep 2022.
- [5] Leonardo Angelini, Massimo Mecella, Hai-Ning Liang, Maurizio Caon, Elena Mugellini, Omar Abou Khaled, and Danilo Bernardini. 2022. Towards an Emotionally Augmented Metaverse: A Framework for Recording and Analysing Physiological Data and User Behaviour. In *13th Augmented Human International Conference (Winnipeg, MB, Canada) (AH2022)*. Association for Computing Machinery, New York, NY, USA, Article 3, 5 pages. <https://doi.org/10.1145/3532530.3532546>
- [6] Apple. 2023. Vision Pro. <https://www.apple.com/apple-vision-pro/>. Online; accessed 6 August 2023.
- [7] Abdullah Al Arafat, Zhishan Guo, and Amro Awad. 2021. VR-Spy: A Side-Channel Attack on Virtual Key-Logging in VR Headsets. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 564–572. <https://doi.org/10.1109/VR50410.2021.00081>
- [8] Ronald Azuma. 1997. A Survey of Augmented Reality. *Presence: Teleoper. Virtual Environ.* 6, 4 (1997). <https://doi.org/10.1162/pres.1997.6.4.355>
- [9] Ronald Azuma, Yohan Baillet, Reinhold Behringer, Steven Feiner, Simon Julier, and Blair MacIntyre. 2001. Recent advances in augmented reality. *IEEE Computer Graphics and Applications* 21, 6 (2001). <https://doi.org/10.1109/38.963459>
- [10] Stefano Baldassi, Tadayoshi Kohno, Franziska Roesner, and Moqian Tian. 2018. Challenges and New Directions in Augmented Reality, Computer Security, and Neuroscience – Part 1: Risks to Sensation and Perception. <https://doi.org/10.48550/ARXIV.1806.10557>
- [11] Matthew Ball. 2022. *The Metaverse: And How It Will Revolutionize Everything*. Minneapolis: Norton & Company.
- [12] Avi Bar-Zeev. 2022. Meta Quest Pro falls short on biometric protections. <https://uxdesign.cc/meta-quest-pro-falls-short-on-biometric-protections-ba48db35637f>. Online; accessed 7 November 2022.
- [13] Steve Benford, Chris Brown, Gail Reynard, and Chris Greenhalgh. 1996. Shared Spaces: Transportation, Artificiality, and Spatiality. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW '96)*. Association for Computing Machinery. <https://doi.org/10.1145/240080.240196>
- [14] Guillermo Bernal, Nelson Hidalgo, Conor Russomanno, and Pattie Maes. 2022. Galea: A physiological sensing system for behavioral research in Virtual Environments. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 66–76. <https://doi.org/10.1109/VR51125.2022.00024>
- [15] bHaptics. 2022. TactSuit X40. <https://www.bhaptics.com/tactsuit/tactsuit-x40>. Online; accessed 22 Sep 2022.
- [16] Courtney Bowman, Ari Gesher, John K. Grant, and Daniel Slate. 2015. *The Architecture of Privacy*. O'Reilly.
- [17] Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F. Schaefer, and Enkelejda Kasneci. 2021. Differential privacy for eye tracking with temporal correlations. *PLOS ONE* 16, 8 (2021), e0255979. <https://doi.org/10.1371/journal.pone.0255979>
- [18] Bracket Foundation. 2022. Gaming and the Metaverse: The Alarming Rise of Online Sexual Exploitation and Abuse of Children Within the NEw Digital Frontier. https://www.weprotect.org/wp-content/uploads/Gaming_and_the_Metaverse_Report_final.pdf. Online; accessed 10 Oct 2022.
- [19] Karla Brkic, Ivan Sikiric, Tomislav Hrkac, and Zoran Kalafatic. 2017. I Know That Person: Generative Full Body and Face De-identification of People in Images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 1319–1328. <https://doi.org/10.1109/CVPRW.2017.173>
- [20] Yavuz Canbay, Anil Utku, and Pelin Canbay. 2022. Privacy Concerns and Measures in Metaverse: A Review. In *2022 15th International Conference on Information Security and Cryptography (ISCTURKEY)*. 80–85. <https://doi.org/10.1109/ISCTURKEY56345.2022.9931866>
- [21] Peter Casey, Ibrahim Baggili, and Ananya Yarramreddy. 2021. Immersive Virtual Reality Attacks and the Human Joystick. *IEEE Transactions on Dependable and Secure Computing* 18, 2 (2021), 550–562. <https://doi.org/10.1109/TDSC.2019.2907942>
- [22] Aayush Kumar Chaudhary and Jeff B Pelz. 2020. Privacy-Preserving Eye Videos Using Rubber Sheet Model. In *ACM Symposium on Eye Tracking Research and Applications (Stuttgart, Germany) (ETRA '20 Short Papers)*. Association for Computing Machinery, New York, NY, USA, Article 22, 5 pages. <https://doi.org/10.1145/3379156.3391375>
- [23] Hang Chu, Shugao Ma, Fernando De la Torre, Sanja Fidler, and Yaser Sheikh. 2020. Expressive Telepresence via Modular Codec Avatars. In *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer International Publishing, Cham, 330–345.
- [24] Pietro Cipresso, Irene Alice Chicchi Giglioli, Mariano Alcañiz Raya, and Giuseppe Riva. 2018. The Past, Present, and Future of Virtual and Augmented Reality Research: A Network and Cluster Analysis of the Literature. *Frontiers in Psychology* 9 (2018), 2086. <https://doi.org/10.3389/fpsyg.2018.02086>
- [25] Matthew Crain. 2021. *Profit Over Privacy*. Minneapolis: University of Minnesota Press.
- [26] Loris D'Antoni, Alan Dunn, Suman Jana, Tadayoshi Kohno, Benjamin Livshits, David Molnar, Alexander Moshchuk, Eyal Ofek, Franziska Roesner, Scott Saponas, Margus Veanes, and Helen J. Wang. 2013. Operating System Support for Augmented Reality Applications. In *14th Workshop on Hot Topics in Operating Systems (HotOS XIV)*. USENIX Association, Santa Ana Pueblo, NM.
- [27] Brendan David-John, Kevin Butler, and Eakta Jain. 2022. For Your Eyes Only: Privacy-Preserving Eye-Tracking Datasets. In *2022 Symposium on Eye Tracking Research and Applications (ETRA '22)*. Association for Computing Machinery, Article 10, 6 pages. <https://doi.org/10.1145/3517031.3529618>
- [28] Brendan David-John, Diane Hosfelt, Kevin Butler, and Eakta Jain. 2021. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization & Computer Graphics* 27, 05 (may 2021), 2555–2565. <https://doi.org/10.1109/TVCG.2021.3067787>
- [29] Jaybie A. de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2019. A First Look into Privacy Leakage in 3D Mixed Reality Data. In *Computer Security – ESORICS 2019*, Kazuo Sako, Steve Schneider, and Peter Y. A. Ryan (Eds.). Springer International Publishing, Cham, 149–169.
- [30] Jaybie Agullo de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2019. SafeMR: Privacy-aware Visual Information Protection for Mobile Mixed Reality. In *2019 IEEE 44th Conference on Local Computer Networks (LCN)*. 254–257. <https://doi.org/10.1109/LCN44214.2019.8990850>
- [31] Jaybie A. de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2019. Security and Privacy Approaches in Mixed Reality: A Literature Survey. *ACM Comput. Surv.* 52, 6, Article 110 (oct 2019), 37 pages. <https://doi.org/10.1145/3359626>
- [32] Jaybie A. de Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2020. Conservative Plane Releasing for Spatial Privacy Protection in Mixed Reality. <https://doi.org/10.48550/ARXIV.2004.08029>
- [33] Alexandre Défossez, Charlotte Caucheteux, Jérémy Rapin, Ori Kabeli, and Jean-Rémi King. 2022. Decoding speech from non-invasive brain recordings. <https://doi.org/10.48550/ARXIV.2208.12266>
- [34] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. 2011. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering* 16, 1 (2011), 3–32. <https://doi.org/10.1007/s00766-010-0115-7>
- [35] Roberto Di Pietro and Stefano Cresci. 2021. Metaverse: Security and Privacy Issues. In *2021 Third IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*. 281–288. <https://doi.org/10.1109/TPSISA52974.2021.00032>
- [36] Ellysse Dick. 2021. Balancing User Privacy and Innovation in Augmented and Virtual Reality. *INFORMATION TECHNOLOGY* (2021), 28.
- [37] Nikita Drobyshev, Janya Chelishv, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. 2022. MegaPortraits: One-shot Megapixel Neural Head Avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*.
- [38] Liang Du, Meng Yi, Erik Blasch, and Haibin Ling. 2014. GARP-face: Balancing privacy protection and utility preservation in face de-identification. In *IEEE International Joint Conference on Biometrics*. 1–8. <https://doi.org/10.1109/BTAS.2014.6996249>
- [39] Reyhan Düzgün, Naheem Noah, Peter Mayer, Sanchari Das, and Melanie Volkmann. 2022. SoK: A Systematic Literature Review of Knowledge-Based Authentication on Augmented Reality Head-Mounted Displays. In *Proceedings of the 17th International Conference on Availability, Reliability and Security (Vienna, Austria) (ARES '22)*. Association for Computing Machinery, New York, NY, USA, Article 36, 12 pages. <https://doi.org/10.1145/3538969.3539011>
- [40] Yogesh K. Dwivedi, Laurie Hughes, Abdullah M. Baabdullah, Samuel Ribeiro-Navarrete, Mihalis Giannakis, Mutaz M. Al-Debei, and et al. 2022. Metaverse beyond the hype: Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management* 66 (2022), 102542. <https://doi.org/10.1016/j.ijinfomgt.2022.102542>
- [41] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.
- [42] Passant Elagrouty, Mohamed Khamis, Florian Mathis, Diana Irmscher, Ekta Sood, Andreas Bulling, and Albrecht Schmidt. 2023. Impact of Privacy Protection

- Methods of Lifelogs on Remembered Memories. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 508, 10 pages. <https://doi.org/10.1145/3544548.3581565>
- [43] Morgane Evin, Antonio Hidalgo-Munoz, Adolphe James Bôquet, Fabien Moreau, Hélène Tattegrain, Catherine Berthelon, Alexandra Fort, and Christophe Jallais. 2022. Personality trait prediction by machine learning using physiological data and driving behavior. *Machine Learning with Applications* 9 (2022), 100353. <https://doi.org/10.1016/j.mlwa.2022.100353>
- [44] Ben Falchuk, Shoshana Loeb, and Ralph Neff. 2018. The Social Metaverse: Battle for Privacy. *IEEE Technology and Society Magazine* 37, 2 (2018), 52–61. <https://doi.org/10.1109/MTS.2018.2826060>
- [45] Lucas Silva Figueiredo, Benjamin Livshits, David Molnar, and Margus Veanes. 2016. Prepose: Privacy, Security, and Reliability for Gesture-Based Programming. In *2016 IEEE Symposium on Security and Privacy (SP)*. 122–137. <https://doi.org/10.1109/SP.2016.16>
- [46] Batya Friedman and Peter H. Kahn. 2000. New Directions: A Value-Sensitive Design Approach to Augmented Reality. In *Proceedings of DARE 2000 on Designing Augmented Reality Environments* (Elsinore, Denmark) (DARE '00). Association for Computing Machinery, New York, NY, USA, 163a–164. <https://doi.org/10.1145/354666.354694>
- [47] Rod Garratt and Maarten R.C. van Oordt. 2018. Privacy as a Public Good: A Case for Electronic Cash. *Journal of Political Economy* (2018). <https://doi.org/10.1086/714133>
- [48] Gonzalo Munilla Garrido, Johannes Sedlmeir, Ömer Uludag, Ilias Soto Alaoui, Andre Luckow, and Florian Matthes. 2022. Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review. *Journal of Network and Computer Applications* 207 (2022), 103465. <https://doi.org/10.1016/j.jnca.2022.103465>
- [49] Alberto Giarretta. 2022. Security and Privacy in Virtual Reality: A Literature Survey. https://doi.org/10.48009/2_iis_2022_125
- [50] Céline Gressel, Rebekah Overdorf, Inken Hagenstedt, Murat Karaboga, Helmut Lurtz, Michael Raschke, and Andreas Bulling. 2023. Privacy-Aware Eye Tracking: Challenges and Future Directions. *IEEE Pervasive Computing* 22, 1 (2023), 95–102. <https://doi.org/10.1109/MPRV.2022.3228660>
- [51] Aniket Gulhane, Akhil Vyas, Reshmi Mitra, Roland Oruche, Gabriela Hofer, Samaikya Valluripally, Prasad Calyam, and Khaza Anuarul Hoque. 2019. Security, Privacy and Safety Risk Assessment for Virtual Reality Learning Environment Applications. In *2019 16th IEEE Annual Consumer Communications Networking Conference (CCNC)*. 1–9. <https://doi.org/10.1109/CCNC.2019.8651847>
- [52] Rajesh Gupta, Sudeep Tanwar, Sudhanshu Tyagi, and Neeraj Kumar. 2020. Machine Learning Models for Secure Data Analytics: A taxonomy and threat model. *Computer Communications* 153 (2020), 406–440. <https://doi.org/10.1016/j.comcom.2020.02.008>
- [53] Jaybie Agullo de Guzman, Aruna Seneviratne, and Kanchana Thilakarathna. 2021. Unravelling Spatial Privacy Risks of Mobile Mixed Reality Data. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 1, Article 14 (mar 2021), 26 pages. <https://doi.org/10.1145/3448103>
- [54] haptx. 2022. Haptx Gloves DK2. <https://haptx.com/>. Online; accessed 22 Sep 2022.
- [55] Ragib Hasan, Suvda Myagmar, Adam J. Lee, and William Yurcik. 2005. Toward a Threat Model for Storage Systems. In *Proceedings of the 2005 ACM Workshop on Storage Security and Survivability*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/1103780.1103795>
- [56] Olli I. Heimo, Kai K. Kimppa, Seppo Helle, Timo Korkalainen, and Teijo Lehtonen. 2014. Augmented reality - Towards an ethical fantasy?. In *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*. 1–7. <https://doi.org/10.1109/ETHICS.2014.6893423>
- [57] VRChat Inc. 2022. The VRChat Security Update. <https://hello.vrchat.com/blog/vrchat-security-update>. Online; accessed 21 October 2022.
- [58] Syem Ishaque, Alice Rueda, Binh Nguyen, Naimul Khan, and Sridhar Krishnan. 2020. Physiological Signal Analysis and Classification of Stress from Virtual Reality Video Game. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. 867–870. <https://doi.org/10.1109/EMBC44109.2020.9176110>
- [59] Suman Jana, David Molnar, Alexander Moshchuk, Alan Dunn, Benjamin Livshits, Helen J. Wang, and Eyal Ofek. 2013. Enabling Fine-Grained Permissions for Augmented Reality Applications with Recognizers. In *22nd USENIX Security Symposium (USENIX Security 13)*. USENIX Association, Washington, D.C., 415–430.
- [60] Suman Jana, Arvind Narayanan, and Vitaly Shmatikov. 2013. A Scanner Darkly: Protecting User Privacy from Perceptual Applications. In *2013 IEEE Symposium on Security and Privacy*. 349–363. <https://doi.org/10.1109/SP.2013.31>
- [61] Carter Jernigan and Behram F.T. Mistree. 2009. Gaydar: Facebook friendships expose sexual orientation. *First Monday* 14, 10 (Sep. 2009). <https://doi.org/10.5210/fm.v14i10.2611>
- [62] Brendan John, Sophie Jörg, Sanjeev Koppal, and Eakta Jain. 2020. The Security-Utility Trade-off for Iris Authentication and Eye Animation for Social Virtual Avatars. *IEEE Transactions on Visualization and Computer Graphics* 26, 5 (2020), 1880–1890. <https://doi.org/10.1109/TVCG.2020.2973052>
- [63] Brendan John, Sanjeev Koppal, and Eakta Jain. 2019. EyeVEIL: Degrading Iris Authentication in Eye Tracking Headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado). New York, NY, USA, Article 37, 5 pages. <https://doi.org/10.1145/3314111.3319816>
- [64] Christos Kalloniatis, Evangelia Kavakli, and Stefanos Gritzalis. 2008. Addressing privacy requirements in system design: the PriS method. *Requirements Engineering* 13, 3 (2008), 241–255. <https://doi.org/10.1007/s00766-008-0067-3>
- [65] Christina Katsini, Yasmeen Abdrabou, George E. Raptis, Mohamed Khamis, and Florian Alt. 2020. The Role of Eye Gaze in Security and Privacy Applications: Survey and Future HCI Research Directions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3313831.3376840>
- [66] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA.
- [67] 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*. Springer International Publishing, Cham, 226–241. https://doi.org/10.1007/978-3-030-42504-3_15
- [68] Jesse Lake. 2020. Hey, You Stole My Avatar!: Virtual Reality and Its Risks to Identity Protection. *EMORY LAW JOURNAL* 69 (2020), 48.
- [69] Hyunjoon Lee, Jiyeon Lee, Daejun Kim, Suman Jana, Insik Shin, and Soeul Son. 2021. AdCube: WebVR Ad Fraud and Practical Confinement of Third-Party Ads. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 2543–2560.
- [70] Sarah M. Lehman, Abrar S. Alrumayh, Kunal Kolhe, Haibin Ling, and Chiu C. Tan. 2022. Hidden in Plain Sight: Exploring Privacy Risks of Mobile Augmented Reality Applications. *ACM Trans. Priv. Secur.* 25, 4, Article 26 (jul 2022), 35 pages. <https://doi.org/10.1145/3524020>
- [71] Jingjie Li, Sunpreet Singh Arora, Kassem Fawaz, Younghyun Kim, Can Liu, Sebastian Meiser, Mohsen Minaei, Maliheh Shirvanian, and Kim Wagner. 2023. "I Want the Payment Process to be Cool": Understanding How Interaction Factors into Security and Privacy Perception of Authentication in Virtual Reality. arXiv:2303.11575 [cs.CR]
- [72] Jingjie Li, Amrita Roy Chowdhury, Kassem Fawaz, and Younghyun Kim. 2021. Kaleido: Real-Time Privacy Control for Eye-Tracking Systems. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1793–1810.
- [73] Jonathan Liebers, Mark Abdelaziz, Lukas Mecke, Alia Saad, Jonas Auda, Uwe Gruenefeld, Florian Alt, and Stefan Schneegass. 2021. Understanding User Identification in Virtual Reality Through Behavioral Biometrics and the Effect of Body Normalization. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 517, 11 pages. <https://doi.org/10.1145/3411764.3445528>
- [74] Zhen Ling, Zupei Li, Chen Chen, Junzhou Luo, Wei Yu, and Xinwen Fu. 2019. I Know What You Enter on Gear VR. In *2019 IEEE Conference on Communications and Network Security (CNS)*. 241–249. <https://doi.org/10.1109/CNS.2019.8802674>
- [75] Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. 2019. Differential Privacy for Eye-Tracking Data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) (ETRA '19). Association for Computing Machinery, New York, NY, USA, Article 28, 10 pages. <https://doi.org/10.1145/3314111.3319823>
- [76] Shiqing Luo, Anh Nguyen, Chen Song, Feng Lin, Wenyao Xu, and Zhisheng Yan. 2020. OcuLock: Exploring Human Visual System for Authentication in Virtual Reality Head-mounted Display. In *NDS5*.
- [77] Magic Leap, Inc. 2022. Magic Leap 2. <https://www.magicleap.com/magic-leap-2>. Online; accessed 4 Oct 2022.
- [78] Divine Maloney, Samaneh Zamanifard, and Guo Freeman. 2020. Anonymity vs. Familiarity: Self-Disclosure and Privacy in Social Virtual Reality. In *26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) (VRST '20). Association for Computing Machinery, New York, NY, USA, Article 25, 9 pages. <https://doi.org/10.1145/3385956.3418967>
- [79] Ivan Martinovic, Doug Davies, Mario Frank, Daniele Perito, Tomas Ros, and Dawn Song. 2012. On the Feasibility of Side-Channel Attacks with Brain-Computer Interfaces. In *21st USENIX Security Symposium (USENIX Security 12)*. USENIX Association, Bellevue, WA, 143–158.
- [80] Ifigenia Mavridou, Ellen Seiss, Theodoros Kostoulas, Charles Nduka, and Emili Balaguer-Ballester. 2018. Towards an Effective Arousal Detection System for Virtual Reality. In *Proceedings of the Workshop on Human-Habitat for Health (H3): Human-Habitat Multimodal Interaction for Promoting Health and Well-Being in the Internet of Things Era* (Boulder, Colorado) (H3 '18). Association

- for Computing Machinery, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/3279963.3279969>
- [81] MelonLoader community. 2022. Melon Loader. <https://melonwiki.xyz>. Online; accessed 22 July 2022.
- [82] Meta. 2022. Eye Tracking Privacy Notice. <https://www.meta.com/en-gb/help/quest/articles/accounts/privacy-information-and-settings/eye-tracking-privacy-notice/>. Online; accessed 21 October 2022.
- [83] Meta. 2022. Meta Quest 2. <https://www.meta.com/es/en/quest/>. Online; accessed 22 Sep 2022.
- [84] Meta. 2022. Meta Quest Pro. https://www.oculus.com/blog/meta-quest-pro-price-release-date/?intern_source=blog&intern_content=meta-quest-pro-privacy. Online; accessed 21 October 2022.
- [85] Meta. 2022. Meta Quest Pro: Built With Privacy in Mind. <https://www.oculus.com/blog/meta-quest-pro-privacy/>. Online; accessed 21 October 2022.
- [86] Meta. 2023. Multitasking, Accessibility Improvements, and Air Link for Quest 1 in the Latest Oculus Software Update. <https://www.meta.com/en-gb/blog/quest/multitasking-accessibility-improvements-and-air-link-for-quest-1-in-the-latest-oculus-software-update/>. Online; accessed 18 May 2023.
- [87] Microsoft. 2022. Azure Automated Machine Learning - AutoML | Microsoft Azure. <https://azure.microsoft.com/en-us/services/machine-learning/automatedml/>
- [88] Stuart E. Middleton, David C. De Roure, and Nigel R. Shadbolt. 2001. Capturing Knowledge of User Preferences: Ontologies in Recommender Systems. In *Proceedings of the 1st International Conference on Knowledge Capture (Victoria, British Columbia, Canada) (K-CAP '01)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/500737.500755>
- [89] Paul Milgram and Fumio Kishino. 1994. A Taxonomy of Mixed Reality Visual Displays. *IEICE Transactions on Information and Systems* 77, 12 (1994), 1321–1329.
- [90] Mark Roman Miller, Fernanda Herrera, Hanseul Jun, James A. Landay, and Jeremy N. Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10, 1 (2020), 17404. <https://doi.org/10.1038/s41598-020-74486-y>
- [91] Robert Miller, Natasha Kholgade Banerjee, and Sean Banerjee. 2022. Combining Real-World Constraints on User Behavior with Deep Neural Networks for Virtual Reality (VR) Biometrics. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 409–418. <https://doi.org/10.1109/VR51125.2022.00060>
- [92] Vivek Nair, Gonzalo Munilla Garrido, and Dawn Song. 2023. Exploring the Privacy Risks of Adversarial VR Game Design. <https://petsymposium.org/2023/files/papers/issue4/popets-2023-0108.pdf>
- [93] Vivek Nair, Gonzalo Munilla Garrido, and Dawn Song. 2023. Going Incognito in the Metaverse: Achieving Theoretically Optimal Privacy-Usability Tradeoffs in VR.
- [94] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data.
- [95] Vivek Nair, Christian Rack, Wenbo Guo, Rui Wang, Shuixian Li, Brandon Huang, Atticus Cull, James F. O'Brien, Marc Latoschik, Louis Rosenberg, and Dawn Song. 2023. Inferring Private Personal Attributes of Virtual Reality Users from Head and Hand Motion Data. [arXiv:2305.19198 \[cs.HC\]](https://arxiv.org/abs/2305.19198)
- [96] Nair, Vivek and Munilla Garrido, Gonzalo. 2022. MetaData Study. <https://github.com/MetaGuard/MetaData>. Online; accessed 27 September 2022.
- [97] Nair, Vivek and Munilla Garrido, Gonzalo. 2022. MetaGuard Study. <https://github.com/MetaGuard>. Online; accessed 27 September 2022.
- [98] Andreas Nautsch, Abelino Jiménez, Amos Treiber, Jascha Kolberg, Catherine Jasserand, Els Kindt, Héctor Delgado, Massimiliano Todisco, Mohamed Amine Hmani, Aymen Mtibaa, Mohammed Ahmed Abdelraheem, Alberto Abad, Francisco Teixeira, Driss Matrouf, Marta Gomez-Barrero, Dijana Petrovska-Delacrétaz, Gérard Chollet, Nicholas Evans, Thomas Schneider, Jean-Francois Bonastre, Bhiksha Raj, Isabel Trancoso, and Christoph Busch. 2019. Preserving privacy in speaker and speech characterisation. *Computer Speech & Language* 58 (2019), 441–480. <https://doi.org/10.1016/j.csl.2019.06.001>
- [99] Neurospec. 2022. DSI-VR300. <https://wearablesensing.com/dsi-vr300/>. Online; accessed 22 Sep 2022.
- [100] Naheem Noah, Sommer Shearer, and Sanchari Das. 2020. Security and Privacy Evaluation of Popular Augmented and Virtual Reality Technologies. In *In Proceedings of the 2022 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence, and Neural Engineering (IEEE MetroXR/AINe 2022)*. Association for Computing Machinery. <https://doi.org/10.2139/ssrn.4173372>
- [101] Fiachra O'Brolcháin, Tim Jacquemard, David Monaghan, Noel O'Connor, Peter Novitzky, and Bert Gordijn. 2016. The Convergence of Virtual Reality and Social Networks: Threats to Privacy and Autonomy. *Science and Engineering Ethics* 22, 1 (2016), 1–29. <https://doi.org/10.1007/s11948-014-9621-1>
- [102] Blessing Odeleye, George Loukas, Ryan Heartfield, Georgia Sakellari, Emanouil Panaousis, and Fotios Spyridonis. 2022. Virtually Secure: A taxonomic assessment of cybersecurity challenges in virtual reality environments. *Computers & Security* (2022), 102951. <https://doi.org/10.1016/j.cose.2022.102951>
- [103] Blessing Odeleye, George Loukas, Fotios Spyridonis, and Ryan Heartfield. 2021. Detecting framerate-oriented cyber attacks on user experience in virtual reality. *USENIX Symposium on Usable Privacy and Security (SOUPS)* (2021), 5.
- [104] UK's Information Commissioner's Office. 2020. Audits of data protection compliance by UK political parties. <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2020/11/uk-political-parties-must-improve-data-protection-practices/>. Online; accessed 17 May 2022.
- [105] Ilesanni Olade, Charles Fleming, and Hai-Ning Liang. 2020. BioMove: Biometric User Identification from Human Kinesiological Movements for Virtual Reality Systems. *Sensors* 20, 10 (2020). <https://doi.org/10.3390/s20102944>
- [106] Seiya Otsuka, Kanami Kurosaki, and Mitsuhiro Ogawa. 2017. Physiological measurements on a gaming virtual reality headset using photoplethysmography: A preliminary attempt at incorporating physiological measurement with gaming. In *TENCON 2017 - 2017 IEEE Region 10 Conference*. 1251–1256. <https://doi.org/10.1109/TENCON.2017.8228049>
- [107] OVR Technology. 2022. Scent Technology for Virtual Reality. <https://ovrtechnology.com/>. Online; accessed 22 Sep 2022.
- [108] Marcus Pendleton, Richard Garcia-Lebron, Jin-Hee Cho, and Shouhuai Xu. 2016. A Survey on Systems Security Metrics. *ACM Comput. Surv.* 49, 4, Article 62 (dec 2016), 35 pages. <https://doi.org/10.1145/3005714>
- [109] Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. 2007. Survey of Network-Based Defense Mechanisms Countering the DoS and DDoS Problems. *ACM Comput. Surv.* 39, 1 (2007). <https://doi.org/10.1145/1216370.1216373>
- [110] Marco Pennacchiotti and Ana-Maria Popescu. 2021. A Machine Learning Approach to Twitter User Classification. *Proceedings of the International AAAI Conference on Web and Social Media* 5, 1 (Aug. 2021), 281–288.
- [111] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural Biometrics in VR: Identifying People from Body Motion and Relations in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 14–21. <https://doi.org/10.1145/3290605.3300340>
- [112] PlayStation. 2022. PlayStation VR2. <https://www.playstation.com/en-us/ps-vr2/>. Online; accessed 28 September 2022.
- [113] CORE Conference Portal. 2023. CORE Computing Research & Education. <http://portal.core.edu.au/conf-ranks/>. Online; accessed 5 August 2023.
- [114] Alan Dexter published. 2021. Oculus will sell you a Quest 2 headset that doesn't need Facebook for an extra \$500. *PC Gamer* (April 2021). <https://www.pcgamer.com/oculus-will-sell-you-a-quest-2-headset-that-doesnt-need-facebook-for-an-extra-dollar500/>
- [115] Ihsan Rabbi and Sehat Ullah. 2016. A Survey on Augmented Reality Challenges and Tracking. *Acta Graphica* 24 (2016), 29–46.
- [116] Muhammad Usman Rafique and Sen-ching S. Cheung. 2020. Tracking Attacks on Virtual Reality Systems. *IEEE Consumer Electronics Magazine* 9, 2 (2020), 41–46. <https://doi.org/10.1109/MCE.2019.2953741>
- [117] Derek Reilly, Mohamad Salimian, Bonnie MacKay, Niels Mathiasen, W. Keith Edwards, and Juliano Franz. 2014. SecSpace: Prototyping Usable Privacy and Security for Mixed Reality Collaborative Environments. In *Proceedings of the 2014 ACM SIGCHI Symposium on Engineering Interactive Computing Systems (Rome, Italy) (EICS '14)*. Association for Computing Machinery, New York, NY, USA, 273–282. <https://doi.org/10.1145/2607023.2607039>
- [118] Black Rock. 2023. The metaverse: Investing in the future now. <https://www.blackrock.com/us/individual/insights/metaverse-investing-in-the-future>. (2023). Online; accessed 7 May 2023.
- [119] Rafael A. Rodríguez-Gómez, Gabriel Maciá-Fernández, and Pedro García-Teodoro. 2013. Survey and Taxonomy of Botnet Research through Life-Cycle. *ACM Comput. Surv.* 45, 4, Article 45 (aug 2013), 33 pages. <https://doi.org/10.1145/2501654.2501659>
- [120] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and Privacy for Augmented Reality Systems. *Commun. ACM* 57, 4 (apr 2014), 88–96. <https://doi.org/10.1145/2580723.2580730>
- [121] Louis B. Rosenberg. 2022. Regulating the Metaverse, a Blueprint for the Future. In *Extended Reality*, Lucio Tommaso De Paolis, Pasquale Arpaia, and Marco Sacco (Eds.). Springer International Publishing, Cham, 263–272.
- [122] Justas Salkevičius, Robertas Damasevičius, Rytis Maskeliūnas, and Ilona Laukienė. 2019. Anxiety Level Recognition for Virtual Reality Therapy System Using Physiological Signals. *Electronics* 8, 9 (2019). <https://doi.org/10.3390/electronics8091039>
- [123] Chamara Sandeepa, Shen Wang, and Madhusanka Liyanage. 2023. Privacy of the Metaverse: Current Issues, AI Attacks, and Possible Solutions. *IEEE International Conference on Metaverse Computing, Networking and Applications (IEEE MetaCom 2023)*.
- [124] Daniel Schneider, Alexander Otte, Travis Gesslein, Philipp Gagel, Bastian Kuth, Mohamad Shahm Damlakhi, Oliver Dietz, Eyal Ofek, Michel Pahud, Per Ola Kristensson, Jörg Müller, and Jens Grubert. 2019. ReconViguration: Reconfiguring Physical Keyboards in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 25, 11 (2019), 3190–3201. <https://doi.org/10.1109/TVCG.>

- 2019.2932239
- [125] Yiran Shen, Hongkai Wen, Chengwen Luo, Weitao Xu, Tao Zhang, Wen Hu, and Daniela Rus. 2019. GaitLock: Protect Virtual and Augmented Reality Headsets Using Gait. *IEEE Transactions on Dependable and Secure Computing* 16, 3 (2019), 484–497. <https://doi.org/10.1109/TDSC.2018.2800048>
- [126] Cong Shi, Xiangyu Xu, Tianfang Zhang, Payton Walker, Yi Wu, Jian Liu, Nitesh Saxena, Yingying Chen, and Jiadi Yu. 2021. Face-Mic: Inferring Live Speech and Speaker Identity via Subtle Facial Dynamics Captured by AR/VR Motion Sensors. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (New Orleans, Louisiana) (*MobiCom '21*). Association for Computing Machinery, New York, NY, USA, 478â&A\$490. <https://doi.org/10.1145/3447993.3483272>
- [127] Shiftall. 2022. HaritoraX 1.1. <https://en.shiftall.net/products/haritorax>. Online; accessed 22 Sep 2022.
- [128] Prakash Shrestha and Nitesh Saxena. 2017. An Offensive and Defensive Exposition of Wearable Computing. *ACM Comput. Surv.* 50, 6, Article 92 (nov 2017), 39 pages. <https://doi.org/10.1145/3133837>
- [129] SJR. 2023. Scimago Journal & Country Rank. <https://www.scimagojr.com/journalsearch.php?q=23038&tip=sid&clean=0>. Online; accessed 5 Aug 2023.
- [130] Carter Slocum, Yicheng Zhang, Jiasi Chen, and Nael Abu-Ghazaleh. 2023. Going through the motions: AR/VR keylogging from user head motions. In *USENIX Security Symposium*.
- [131] Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-Aware Eye Tracking Using Differential Privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) (*ETRA '19*). Association for Computing Machinery, New York, NY, USA, Article 27, 9 pages. <https://doi.org/10.1145/3314111.3319915>
- [132] Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-Aware Eye Tracking Using Differential Privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) (*ETRA '19*). Association for Computing Machinery, New York, NY, USA, Article 27, 9 pages. <https://doi.org/10.1145/3314111.3319915>
- [133] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019. PrivacEye: Privacy-Preserving Head-Mounted Eye Tracking Using Ego-centric Scene Image and Eye Movement Features. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado) (*ETRA '19*). Association for Computing Machinery, New York, NY, USA, Article 26, 10 pages. <https://doi.org/10.1145/3314111.3319913>
- [134] Sophie Stephenson, Bijeta Pal, Stephen Fan, Earlene Fernandes, Yuhang Zhao, and Rahul Chatterjee. 2022. SoK: Authentication in Augmented and Virtual Reality. In *2022 IEEE Symposium on Security and Privacy (SP)*. 267–284. <https://doi.org/10.1109/SP46214.2022.9833742>
- [135] Philipp Sykownik, Divine Maloney, Guo Freeman, and Maic Masuch. 2022. Something Personal from the Metaverse: Goals, Topics, and Contextual Factors of Self-Disclosure in Commercial Social VR. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 632, 17 pages. <https://doi.org/10.1145/3491102.3502008>
- [136] Piotr Szczuko. 2014. Augmented Reality for Privacy-Sensitive Visual Monitoring. In *Multimedia Communications, Services and Security*, Andrzej Dziech and Andrzej Czyżewski (Eds.). Springer International Publishing, Cham, 229–241.
- [137] Luma Tabbaa, Ryan Searle, Saber Mirzaee Bafti, Md Moinul Hossain, Jitrapol Intarasrisrisawat, Maxine Glancy, and Chee Siang Ang. 2022. VREED: Virtual Reality Emotion Recognition Dataset Using Eye Tracking & Physiological Measures. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4, Article 178 (dec 2022), 20 pages.
- [138] Jerry Tang, Amanda LeBel, Shailee Jain, and Alexander G. Huth. 2022. Semantic reconstruction of continuous language from non-invasive brain recordings. *bioRxiv* (2022). <https://doi.org/10.1101/2022.09.29.509744>
- [139] Robert Templeman, Mohammed Korayem, David Crandall, and Apu Kapadia. 2014. PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. In *Proceedings 2014 Network and Distributed System Security Symposium* (San Diego, CA). Internet Society. <https://doi.org/10.14722/ndss.2014.23014>
- [140] Financial Times. 2022. Facebook Patents Reveal How It Intends to Cash In On Metaverse. <https://www.ft.com/content/76d40aac-034e-4e0b-95eb-c5d34146f647>. Online; accessed 21 October 2022.
- [141] Andrew Trask, Emma Bluemke, Ben Garfinkel, Claudia Ghezou Cuervas-Mons, and Allan Dafoe. 2020. Beyond Privacy Trade-offs with Structured Transparency. *arXiv:2012.08347 [cs.CR]*
- [142] Pier Paolo Tricomi, Federica Nenna, Luca Pajola, Mauro Conti, and Luciano Gamberini. 2023. You Can't Hide Behind Your Headset: User Profiling in Augmented and Virtual Reality. *IEEE Access* 11 (2023), 9859–9875. <https://doi.org/10.1109/ACCESS.2023.3240071>
- [143] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 3789–3806.
- [144] Wen-Jie Tseng, Elise Bonnal, Mark McGill, Mohamed Khamis, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. 2022. The Dark Side of Perceptual Manipulations in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 612, 15 pages. <https://doi.org/10.1145/3491102.3517728>
- [145] Vive. 2022. Facial Tracker. <https://www.vive.com/us/accessory/facial-tracker/>. Online; accessed 4 Oct 2022.
- [146] VIVE. 2022. Introducing VIVE Tracker. <https://www.vive.com/us/accessory/tracker3/>. Online; accessed 21 Sep 2022.
- [147] Vive. 2022. Vive Cosmos. <https://www.vive.com/us/product/vive-cosmos/overview/>. Online; accessed 22 Sep 2022.
- [148] Vive. 2022. Vive Pro. <https://www.vive.com/us/product/vive-pro-full-kit/>. Online; accessed 22 Sep 2022.
- [149] Vive. 2022. Vive Pro Eye. <https://www.vive.com/us/product/vive-pro-eye/overview/>. Online; accessed 22 Sep 2022.
- [150] Martin Vondráček, Ibrahim Baggili, Peter Casey, and Mehdi Mekni. 2022. Rise of the Metaverse's Immersive Virtual Reality Malware and the Man-in-the-Room Attack & Defenses. *Computers & Security* (2022), 102923. <https://doi.org/10.1016/j.cose.2022.102923>
- [151] VRChat. 2022. Network Specs and Tips. <https://docs.vrchat.com/docs/network-details>. Online; accessed 4 Oct 2022.
- [152] Junjue Wang, Brandon Amos, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan. 2017. A Scalable and Privacy-Aware IoT Service for Live Video Analytics. In *Proceedings of the 8th ACM on Multimedia Systems Conference* (Taipei, Taiwan) (*MMSys'17*). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3083187.3083192>
- [153] Yuntao Wang, Zhou Su, Ning Zhang, Rui Xing, Dongxiao Liu, Tom H. Luan, and Xuemin Shen. 2022. A Survey on Metaverse: Fundamentals, Security, and Privacy. *IEEE Communications Surveys & Tutorials* (2022), 1–1. <https://doi.org/10.1109/COMST.2022.3202047>
- [154] Xing Wei and Chenyang Yang. 2022. FoV Privacy-aware VR Streaming. In *2022 IEEE Wireless Communications and Networking Conference (WCNC)*. 1515–1520. <https://doi.org/10.1109/WCNC51071.2022.9771832>
- [155] Yu-Chih Wei, Wei-Chen Wu, Gu-Hsin Lai, and Ya-Chi Chu. 2020. pISRA: privacy considered information security risk assessment model. *The Journal of Supercomputing* 76, 3 (2020), 1468–1481. <https://doi.org/10.1007/s11227-018-2371-0>
- [156] Alexander Winkler, Jungdam Won, and Yuting Ye. 2022. QuestSim: Human Motion Tracking from Sparse Sensors with Simulated Avatars. In *SIGGRAPH Asia 2022 Conference Papers* (Daegu, Republic of Korea) (*SA '22*). Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. <https://doi.org/10.1145/3550469.3555411>
- [157] Dongrui Wu, Christopher G. Courtney, Brent J. Lance, Shrikanth S. Narayanan, Michael E. Dawson, Kelvin S. Oie, and Thomas D. Parsons. 2010. Optimal Arousal Identification and Classification for Affective Computing Using Physiological Signals: Virtual Reality Stroop Task. *IEEE Transactions on Affective Computing* 1, 2 (2010), 109–118. <https://doi.org/10.1109/T-AFFC.2010.12>
- [158] Felix T Wu. 2012. Defining Privacy and Utility in Data Sets. *84 University of Colorado Law Review* 1117 (2013); 2012 *TRPC* (2012), 1117–1177. <https://doi.org/10.2139/ssrn.2031808>
- [159] Yifan Wu, Fan Yang, Yong Xu, and Haibin Ling. 2019. Privacy-Protective-GAN for Privacy Preserving Face De-Identification. *Journal of Computer Science and Technology* 34, 1 (2019), 47–60. <https://doi.org/10.1007/s11390-019-1898-8>
- [160] Eisa Zarepour, Mohammadreza Hosseini, Salil S. Kanhere, and Arcot Sowmya. 2016. A context-based privacy preserving framework for wearable visual lifeloggers. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. 1–4. <https://doi.org/10.1109/PERCOMW.2016.7457057>
- [161] Yicheng Zhang, Carter Slocum, Jiasi Chen, and Nael Abu-Ghazaleh. 2023. It's all in your head(set): Side-channel attacks on AR/VR systems. In *USENIX Security Symposium*.
- [162] Yingying Zhao, Yuhu Chang, Yutian Lu, Yujiang Wang, Mingzhi Dong, Qin Lv, Robert P. Dick, Fan Yang, Tun Lu, Ning Gu, and Li Shang. 2022. Do Smart Glasses Dream of Sentimental Visions? Deep Emotionship Analysis for Eyewear Devices. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 1, Article 38 (mar 2022), 29 pages. <https://doi.org/10.1145/3517250>

A TERMINOLOGY




Figure 4: Milgram and Kishino’s Reality and Virtuality Continuum (cf. adapted [89]).

Augmented Reality (AR) is an interactive environment where virtual content (e.g., text, objects, etc.) is spatially embedded into reality and experienced in the first person [89, 121].

Virtual Reality (VR) is a computer-simulated interactive environment experienced in the first person [89, 121].

Extended Reality (XR) and **Mixed Reality (MR)** are umbrella terms that include all forms of immersive media (See Fig. 4). Currently, MR is often used as a synonym for AR [121].

B VIRTUAL REALITY INPUT AND OUTPUT DEVICES

In the following, we compile a succinct list of the VR input and output devices considered in this work and referred to in Fig. 1.

Input devices:

- Head mounted display
 - Face trackers
 - Eye trackers
 - Inside-Out tracking sensors
- Hand-Held controllers
- Gloves
- Full-Body Tracking
- Outside-In Tracking Sensors
- Microphone
- Physical/Virtual Keyboard
- Health sensors

Output devices:

- Graphic display
- Speakers
- Haptic feedback devices
 - Hand-Held controllers
 - Gloves
 - Vests

C DETAILED DATA COLLECTION METHOD & RESULTS

We inspired our method from seminal SoKs and reviews in the field of security and privacy [31, 39, 48]. For this SoK, we sought literature presenting at least one of the following artifacts in the context of virtual reality (VR) and privacy: (i) privacy threat or (ii) defense models, (iii) taxonomy of attributes or (iv) applications, (v) survey or implementation of attacks or (vi) defenses. When a paper encompassed mixed reality, we included the work if the presented artifacts (partially) overlapped with VR.

Before the search, the researchers knew of 12 relevant studies containing the target artifacts (base literature). Two researchers

curated the search string by studying the base literature and conducting a manual preliminary search in Google Scholar for papers containing the targeted artifacts.

Search string: (“virtual reality” OR “virtual telepresence” OR “head-mounted displays” OR “head-worn display” OR “metaverse”) AND “data” AND “privacy” AND (“attack” OR “offense” OR “defense” OR “protection”)

With this search string, we queried on August 10th, 2022 the seven most relevant digital libraries: IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, Scopus, Wiley InterScience, and Web of Science. We included work published in or after 2010 and excluded books, resulting in 1700 hits. Two researchers filtered in parallel the publications by title (47 selected from 1700), abstract (35 selected from 47), and full text (14 selected from 35), resolving conflicts in an informed discussion attending to the target artifacts and removing duplicates. To reduce the number of possible missing publications, we employed Google Scholar to rerun the same search string and process. We stopped when the publication titles did not contain keywords and added two new publications to the selected studies. Combining the base literature (12) with the filtered studies (16) resulted in 23 selected studies after deduplication.

To further ensure we collected as many relevant publications as possible, we conducted a backward search of the references of the 23 selected studies under the same criteria. Additionally, we contacted the authors of the 23 works to obtain further relevant publications. The backward search revealed 26 studies, and from the corpus signaled by the scholars, we included another 7 after deduplication. Lastly, we collected another 19 publications thereafter throughout the research and writing of this manuscript following the same criteria (with publications as recent as May 2023). Table 5 shows our final list of 75 publications, and which of them are not peer-reviewed (only 4). In addition, Table 5 indicates the studies that have been showcased in A* or A publication venues based on the CORE ranking [113] for conferences and SJR [129] for journals (quartile 1). This allows researchers to critically assess the authors’ assertions within this SoK. Remarkably, these high-profile publications comprise 56% of the chosen papers for this review.

D DIFFERENCES OF VR AND AR THROUGH THE LENS OF SECURITY AND PRIVACY

VR and AR are often grouped under the collective term XR or MR. However, from a security and privacy perspective, they have unique characteristics and implications. VR immerses users entirely into a computer-simulated interactive environment, whereas AR spatially embeds virtual content onto the real world through the user’s view [89, 121]. This fundamental distinction drives different interaction paradigms and associated data collection needs.

In VR, the entirety of the user’s visual and auditory sensory input is generated by the device. This total immersion means that the system must capture a comprehensive set of user movements, biometrics, among other data (see §5) to provide a convincing and responsive virtual experience. Such exhaustive data capture naturally introduces a broader surface for potential privacy and security concerns [35]. Furthermore, by shielding the user’s field of view and hearing, VR “removes” external stimuli (e.g., from another individual standing in the same room) making the experience (over-)immersive and

Table 5: The 75 collected studies.

Study focus	Studies (75)
VR Primary Studies (63)	
Defenses (35)	[131], [27], [154], [75], [91] [†] , [93], [28] [‡] , [44], [70] [‡] , [133], [26], [59], [29], [32], [17], [72] [‡] , [22] [‡] , [60], [139], [45], [160], [136], [152], [38], [19], [159], [63], [62], [30] [‡] , [73] [†] , [111] [†] , [125] [†] , [105] [†] , [117], [124]
Attacks (24)	[53], [162] [†] , [3] [‡] , [92], [90], [126], [156] [†] , [150] [‡] , [142], [137] [†] , [80] [†] , [5] [†] , [122] [†] , [106] [†] , [14] [†] , [157] [†] , [58] [†] , [74], [7], [94], [123], [161], [130], [95]
Surveys (2)	[78], [135] [†]
Evaluations (2)	[100], [143]
VR Secondary Studies (12)	
Literature	[31], [40], [35], [65], [36], [128], [153],
Reviews	[67], [49], [102], [50], [20]
Studies in A* and A conferences (32%): [94], [42], [161], [130], [69], [78], [135], [91], [65], [92], [93], [143], [26], [59], [29], [90], [72], [60], [139], [45], [126], [73], [111], [7]	
Studies in quartile 1 journals (24%): [31], [53], [162], [40], [36], [28], [128], [70], [153], [17], [62], [150], [142], [125], [137], [157], [102], [124]	
Selected Studies Not Peer-Reviewed (5%): [32], [49], [71], [95]	

[†]An attacker can leverage the associated defense/mechanism for adversarial purposes. [‡]Although the study is defense/attack focused, there is an adversarial/defensive component.

more vulnerable to deception [40, 92].

On the other hand, AR, while also data-intensive, is anchored in the real world. Its primary data needs revolve around understanding and augmenting the physical environment. This means AR devices continuously scan and interpret the real-world setting, potentially introducing privacy risks related to location tracking, real-world activity monitoring, and the unintentional capture of third-party data [36].

The data needs and unique interaction models of VR and AR mean that certain threats and protections are more pronounced in one domain over the other. For instance, an attack vector exploiting the immersive nature of VR, e.g., deceiving users into sharing sensitive information [92], may not be as effective in AR. Conversely, a privacy concern related to real-world location tracking or user daily habits in AR does not apply to the same extent in VR. However, as we indicated in KF8, there still exists an overlap between AR and VR systems: their frontal cameras and eye trackers, which may be attacked in fundamentally similar ways.

E QUANTITATIVE ANALYSIS RESULTS

Tables 6 and 7 provide the results of our quantitative analysis covered in §8.

Table 6: Cumulative explained variability (%) of PC1 for each primary and secondary attribute per inferred attribute.

Inferred	Primary & Secondary Attributes												
	Demographics	Height	Room Size	Voice	Wingspan	Longer Arm	IPD	Vision	Reaction Time	Game Duration	Location	MoCA Score	Language
Gender	51.07%	-	-	-	33.65%	-	15.28%	-	-	-	-	-	-
Age	26.63%	-	-	-	-	-	-	25.66%	21.44%	17.03%	-	9.24%	-
Ethnicity	14.19%	-	-	84.68%	-	-	-	-	-	-	-	-	1.13%
Income	-	85.76%	-	-	-	-	-	-	-	-	14.24%	-	-
Identity	21.29%	-	-	-	21.37%	37.08%	20.25%	-	-	-	-	-	-
Risk Total:	113.18%	85.76%	84.68%	55.03%	37.08%	35.53%	25.66%	21.44%	17.03%	14.24%	9.24%	1.13%	-

Table 7: Weighted mean degradation of attack accuracy per attribute using MetaGuard [93] as defense mechanism.

Attribute	Accu. at No Priv.	L Priv. Parameter	Accu. at L Priv.	M Priv. Parameter	Accu. at M Priv.	H Priv. Parameter	Accu. at H Priv.	WMD	Risk	Opportunity (Score)
Room Size	97%	$\epsilon=3.00$	33.52%	$\epsilon=1.00$	23.44%	$\epsilon=0.10$	19.53%	66.28%	85.76%	56.84
Height	100%	$\epsilon=5.00$	68.6%	$\epsilon=3.00$	58.17%	$\epsilon=1.00$	44.47%	37.56%	113.18%	42.51
IPD	87%	$\epsilon=5.00$	19.47%	$\epsilon=3.00$	14.17%	$\epsilon=1.00$	12.17%	70.11%	35.53%	24.91
Wingspan	100%	$\epsilon=3.00$	78.80%	$\epsilon=1.00$	66.00%	$\epsilon=0.50$	65.46%	25.53%	55.03%	14.05
Location	90%	$d=25.00$ ms	6.66%	$d=30.00$ ms	0.00%	$d=50.00$ ms	0.00%	88.41%	14.24%	12.59
Voice	63%	$\epsilon=6.00$	52.50%	$\epsilon=1.00$	40.00%	$\epsilon=0.10$	32.50%	12.54%	84.68%	10.62
Longer Arm	100%	$\epsilon=3.00$	77.78%	$\epsilon=1.00$	62.22%	$\epsilon=0.10$	53.33%	26.61%	37.08%	9.87
Reaction Time	87.50%	$d=10.00$ ms	79.20%	$d=20.00$ ms	62.50%	$d=100.00$ ms	54.20%	30.10%	21.44%	6.45
Refresh Rate	100%	$d=90.00$ ms	0.00%	$d=72.00$ ms	0.00%	$d=60.00$ ms	0.00%	100.00%	-	-
Tracking Rate	100%	$d=90.00$ ms	0.00%	$d=72.00$ ms	0.00%	$d=60.00$ ms	0.00%	100.00%	-	-
Handedness	97%	$\epsilon=1.30$	92.50%	$\epsilon=1.00$	75.00%	$\epsilon=0.70$	57.50%	18.50%	-	-
Physical Fitness	90.00%	$\epsilon=5.00$	86.11%	$\epsilon=3.00$	79.11%	$\epsilon=1.00$	61.56%	8.95%	-	-

Abbreviations: d = delay, Accu. = Accuracy, L = Low, M = Medium, H = High, Priv. = Privacy,

$$WMD = \text{Weighted Mean Degradation} = (\text{Accuracy at No Privacy}) - \frac{\sum_{i \in \{L, M, H\}} (\text{Parameter}_i * \text{Accuracy}_i)}{\sum_{i \in \{L, M, H\}} (\text{Parameter}_i)}, \text{Opportunity} = WMD * \text{Risk}.$$